FIRST MAC CLONES

BYTE

THE MAGAZINE OF TECHNOLOGY INTEGRATION

Notes 4.0 PAGE 131

Build Your Own World Wide Web Server

9.50

Microsoft Visual C++ Goes Multiplatform

Intel's P6

In-Depth Analysis

PLUS

30 Best Notebooks

Special Report: The Future of Client/Server Computing
PAGE 105

IMAGINATIONSTATION

When your 486 just doesn't cut it.

THE P75 HOME MPC
Your Best Value in Pentium Power

Face the facts. When it comes to multimedia, your old 486 just can't keep up with the new, faster-than-lightning Pentium Home systems from Micron. By the time your spreadsheet, financial software, word processor and presentation packages are loaded--and the temptation of adding some exciting new educational software and a few hot arcade-quality multimedia games is more than you can resist--your trusty, but dusty 486 cries for help.

Don't despair. . . Micron now offers you an attractive alternative--the Micron P75 Home MPC, with 75MHz of Pentium performance at 486 prices. Plus we include SoundBlaster 16 stereo sound, speakers, CD-ROM drive and fax/modem--plus a wide variety of home productivity and entertainment software on CD. All in one complete package. Why settle for less?

- Intel Pentium 75MHz processor
- 8MB RAM, 256K write-back cache
- 527MB 12ms IDE hard drive (PCI)
- 2X CD-ROM drive, 3.5" floppy drive
- SoundBlaster 16 sound card, speakers & joystick
- 14.4/9600 Fax/Modem & WinFax Lite
- PCI 64-bit graphics accelerator (2MB)
- 14" Micron 14FG, 1024NI, .28mm

$2,299
Complete!

PC WORLD
BEST BUY
January, 1995
P90 Home MPC

Every Micron Home MPC comes free with Microsoft *Encarta* CD, Microsoft *Works Multimedia* CD and *Quicken Deluxe Edition* CD. You also get the latest CD games from LucasArts: *Rebel Assault, Day of the Tentacle, and Indiana Jones and the Fate of Atlantis*. Along with trial subscriptions for CompuServe, America OnLine and Prodigy. All included in the one low affordable system price.

- "Tool-Free" desktop or mini-tower
- Microsoft Mouse, 101-key keyboard
- MS-DOS & Windows for Workgroups
- Microsoft Works Multimedia CD, Microsoft Encarta CD, Quicken Deluxe Edition CD
- Trial subscriptions for Compuserve, America OnLine & Prodigy
- LucasArts CD games: Rebel Assault, Day of the Tentacle, Indiana Jones and the Fate of Atlantis

▶ With 90MHz Pentium Processor...... $2499
▶ With 100MHz Pentium Processor.... $2699

EXCLUSIVELY FROM MICRON
- Next Day Shipping on Selected Systems
- 30-Day Risk-Free Money Back Guarantee
- One-Year Express Parts Replacement
- Lifetime Toll-Free Technical Support
- On-Site Service or System Replacement

MICRON
COMPUTER, INC.
800-233-7027

ENGINEERED
FOR EXCELLENCE

MICRON
COMPUTER, INC.

TAKE MULTIMEDIA INTO

ANNOUNCING THE HIGH-PERFORMANCE T4850CT. Now you can create dazzling multimedia
presentations on a large 10.4" display, with the excitement of Indeo™ video and the richness of integrated
.WAV and MIDI Audio. By connecting this plug-and-play notebook to an optional Desk Station IV,

T4850CT FEATURES:
- 75MHz IntelDX4™, 3.3 volt processor, 16KB cache
- 10.4" dia. color SVGA TFT-LCD
  active matrix display
- 810 Million Byte (= 772MB) HDD or
  520 Million Byte (= 500MB) HDD
- 8MB RAM expandable to 24MB
- Two PCMCIA slots (16mm and 5mm)

- VL local-bus video
- Integrated graphics accelerator
- 65,536 (64K) simultaneous colors
- .WAV Audio and MIDI sound
- Audio jacks: headphone, built-in
  speakers and microphone
- 6.9 lbs.
- NiMH battery for extended life

- 3.5" 1.44MB floppy disk drive
- Ballpoint™ Mouse with QuickPort™
- Pre-installed software: DOS®, Windows
  for Workgroups™, Windows Sound System™,
  Run Time Video for Windows™, Indeo™
  video and Fn-esse™ software
- Toll-free Technical Support –
  7 days a week/24 hours a day

indeo
INTEL VIDEO TECHNOLOGY

Indeo™ video bolstered by a powerful 75MHz IntelDX4™ processor provides greater performance for smoother video playback.

10.4"dia. The 10.4" color active matrix screen displays 65,536 (64K) simultaneous colors for stunningly realistic images and brilliant graphics.

810 MILLION BYTES Choose between a robust 810 Million Byte (= 772MB) or 520 Million Byte (= 500MB) hard disk drive configuration for storing large multimedia files.

Dual PCMCIA expansion slots (Type II and Type III) run simultaneously—with connection to your LAN, fax/modem and more.

Integrated .WAV Audio and MIDI pump out clear and powerful sound effects, music and dialogue.

NEW DIMENSIONS.

you'll have the flexibility to hook up an external monitor, keyboard, mouse and more, for a comfortable desktop environment. The new T4850CT puts the show and the audience in the palm of your hand. **For your nearest Toshiba dealer, call 1-800-457-7777.**

In Touch with Tomorrow
TOSHIBA

Cover Story

Special Report: Client/Server Computing

Features

Reviews

PAGE 165

PAGE 181

CA-V

File   Edit   View   Tools   Application

Applications, Libraries, and DLLs

DBF   DBF Classes

GUI Classes   OE Data Servers

Order Entry   Report Classes

PAGE 179

State of the Art

BYTE Contents by Platform

Index

This page presents the articles in this issue according to computing platform.

YOUR NEXT HARD DRIVE IS A PINNACLE OPTICAL.

Have you outgrown your hard drive? Need more storage? Go optical. There is no doubt that optical is the storage solution for the future. And Pinnacle Micro is the optical storage leader. Optical is fast - fast as most hard drives. Optical is removable - fill up a disk... just add another one. Optical is reliable - no data loss. No head crashes. Even better, optical media costs less than 15 cents per megabyte.

Now you can store unlimited amounts of data for endless applications such as graphics, imaging, digital audio, digital video, CAD/CAM, database, network storage and data archiving. You can even transport Megabytes and Gigabytes of large files within seconds just by handing a disk from one person to another. Or place your disk in a Fed-Ex package and ship it for next day delivery.

Tahoe™ 230 MB
3.5" *Portable Optical Drive*

Sierra™ 1.3 GB
5.25" *Optical Hard Drive*

Pinnacle offers the industry's fastest 3.5" and 5.25" optical drive solutions, the most affordable Recordable CD system, and network optical library systems that range from 20 Gigabytes to over 1 Terabyte for on-line, centralized storage. Pinnacle's award-winning line of optical storage solutions will increase productivity in every environment, from personal computing to the largest global network.

So before you buy your next hard drive... go Pinnacle optical. A storage solution you will never outgrow.

All trademarks and Registered Trademarks of Their Respected Owners

TO REACH THE PINNACLE CALL: 800-553-7070

PINNACLE MICRO
THE OPTICAL STORAGE COMPANY

Circle 91 on Inquiry Card (RESELLERS: 92).

Realtime technology...

Why do so many realtime developers like working with QNX?

POSIX Outside, *Performance* Inside

If you know the UNIX® OS, you know the API and utilities of QNX or any POSIX OS. But when it comes to performance, you can't judge a POSIX OS by its cover – they're all *supposed to* look alike. Check out QNX's numbers (e.g. 5 µsec per context switch on an 80486).

Scale it Down, Up, Out

QNX's modular architecture lets you scale down the OS to fit on a tiny ROM-based platform. Or scale it up to drive a resource-rich workstation equipped with X and TCP/IP.

Or scale it out to build a vast network running hundreds of CPUs. You can even "hot-swap" OS modules (filesystems, device drivers, system processes) *on the fly at run time!*

Tools for *Speed*

You'll like the tighter, faster code you get with the Watcom C/C++ optimizing compiler.

You'll also like our distributed debugger, profiler, trace analysis tools, code-generating GUI builder – they're all engineered to save you time.

Let's Talk *Technology*

If your realtime apps call for a proven, high-performance OS, call us. And ask all the detailed technical questions you can think of. After all, QNX is made by programmers for programmers.

...for the bottom line

Why do so many corporations succeed with QNX?

One OS Fits All

From cash registers to instrumentation, PDAs to robotic controllers, QNX is the only realtime OS that can be scaled to fit the platform. So your programmers won't waste time learning a new OS for every project. And you won't have to make another OS buying decision. Naturally, a single modular OS will cut the cost of your runtime systems, because you pay for only the OS modules you use.

Tools for Productivity

Runtime performance is critical, but so is beating your competitors to market. With QNX's rich development environment, your programmers will have the tools they need to produce better software sooner.

POSIX for Open Systems

As a certified POSIX OS, QNX lets you port applications quickly and easily across platforms. And programmers who know UNIX or any POSIX OS will be productive with QNX immediately.

Let's Talk Business

Why is QNX the leader in realtime OS technology? Proven 14-year track record, satisfied Fortune 500 customers, no-nonsense licensing policy, Gold support plan...

The bottom line? QNX can give your realtime applications a real advantage.

QNX®

The Leading Realtime OS for PCs

- **Leading in Experience** (Realtime OS for PCs since 1981)
- **Leading in Innovation** (Microkernel distributed OS for PCs since 1984)
- **Leading in Market Share** (QNX outsells every other realtime OS for PCs)

800-676-0566 (EXTENSION 1004)

Circle 95 on Inquiry Card.

The In-Your-Face Interface

Why Bob matters. The good, the bad, and the cranky.

Ask almost any computer expert their opinion of Bob, Microsoft's new interface/application suite with "personality," and you'll get a cringe, a wince, a wrinkle of the nose. No thanks, too cute, too slow, no APIs. Melinda French of Microsoft even said recently to a roomful of 800 computer-industry movers and shakers, "this is a program that nobody in this audience will use."

But Bob matters, and not just because computer wonks like you and me will be spending more time troubleshooting our in-laws' new triple-layer-cake systems (Bob is a shell on top of Windows, which runs on DOS). No, Bob matters because it has several crucial user-interface innovations that will change the way the next generation of computer users expect their machines to behave. Bob's innovations are going to make systems design more complicated.

My Friend Scuzz

The most visible interface element of Bob is its little cartoon-character personalities that act as your guides and are the center of Bob's "social interface." My favorite is Scuzz, the adolescent rat, who comes with props like a basketball and an out-of-tune electric guitar. Microsoft people will tell you they based the concept of Bob on a well-known fact: Users often react to their PCs as if they are living beings—they smile, get angry, and talk to them. So, based on the fact that people treat their computers as if they are alive and have personality, Microsoft designed Bob so the personalities are in your face.

Is this a good thing? Like a lot of people, I talk to my car, mostly when it's not doing what I want it to. But it's just an inanimate hunk of metal and plastic. Its personality is the one I project onto it, affected by the design decisions of its engineers, especially the one who designed in the loud clunk it emits when accelerating out of a corner. I am happy—delighted, actually—that the same engineer didn't design my car with an overt personality. As it is, it has all the personality I can handle.

The problem is, users who get used to Bob and Bob-like interfaces are going to start expecting new programs to have overt personalities. When your sister goes to get money out of Mister ATM, she'll expect it to have a Mister attitude. In the absence of a character-based helping hand, people are going to feel lost.

I don't want to sound like a Luddite, because I'm in favor of technology that makes computers easier to use. But there's a programming problem here for systems implementers. What personality will new custom programs have? If you write a system for your company, how will you design its personality (as if its bugs and quirks aren't personality enough)? Will you have to hire a personality consultant for every new project? Or will some clever vendor create a toolkit of personalities you can plug into your application? Regardless, you'd better start thinking about it, because users raised on Bob are going to expect something—little smurfy bunnies, obnoxious rats, or other jabbering cartoons—to help walk them through new applications.

Do What I Mean

The other key tenet of Bob is that its personalities track what you do. Do the same operation an inefficient way enough times, and the cartoon will pop up a little text balloon that says something like, "Are you trying to erase a whole word? In the future, you can press Shift-Backspace to do the same thing quicker."

Now, aside from the fact that I don't normally take advice from cartoon characters, this concept is spot-on. Computer products should anticipate the needs of the user. Unfortunately, few programmers have enough training in human-factors engineering to understand all the misguided paths a user might take to perform a given task. I've sat in on plenty of usability tests where it became obvious why the video booth is soundproofed: It's not to keep the sound of the person under test from getting into the control room; it's to keep the programmers' screams and yelps of dismay from distracting the user.

It's hard to write simple programs. It takes a special sensitivity and an ability to remove yourself from your project long enough to see it through the eyes of users. It also takes a lot of resources and code to make the core of any program approachable from the thousand different directions from which a thousand users will come at it.

But save me from cartoons. Give me a button labeled "Do What I Mean," and leave it at that, thank you very much. ■

RAPHAEL NEEDLEMAN, EDITOR IN CHIEF
(rafe@well.com)

Olivetti Echos 20 Series
Before choosing an ultralight portable PC, count to ten.

1 THE 1.8-KG COLOUR PORTABLE

Echos 20/C is a truly ultralight colour portable, weighing in at only 1.8 kg. Its compact dimensions are best expressed in millimeters: 264 x 190 x 38. Just wide enough not to skimp on the keyboard, just deep enough to mount the biggest screen in its class, and just thick enough to contain all the features you need. And Echos 20 with a monochromatic screen is even slimmer - a mere 35 mm, or 1 3/16 inches.

4 9.4 -INCH DSTN COLOUR DISPLAY

Some things just cannot be miniaturised (like your presentations), so Echos 20/C has a new generation high-contrast 9.4" DSTN colour display. This excellent screen is as big as any in the subnotebook class, and bigger than many found on heavier and costlier portables. A Vesa local bus and new Fast IDE interface interact to elaborate graphics in 256 colours with surprising speed.

2 HEAVYWEIGHT PERFORMANCE

This ruggedly elegant little package outperforms many bulkier machines. Your graphics really move thanks to a Fast IDE interface and Vesa local bus, and Echos' 4MB of RAM easily expands to 20 MB. With HDU ranging from 170 to 240 MB and top-performing Intel 486 CPUs, Echos is a no-compromise solution for your business.

3 LOW CONSUMPTION 3.3-VOLT CIRCUITRY

Unlike most portables, Olivetti Echos employs energy efficient 3.3-Volt circuitry for its CPU, chipset and memory. That means Echos 20 with B/W screen can sustain five consecutive hours of intensive use (run-down test with NiMH battery). Its APM-compliant power management system is so efficient that it can stay in suspend mode for 143 hours - almost 6 days.

5 PLUG AND PLAY EASE

A Phoenix plug and play driver ensures that Echos instantly recognises a vast array of PCMCIA cards. Olivetti

PCMCIA type III slot

has also equipped these ultralights with print drivers for more than 50 major types of printers, so you can print easily using almost anyone's printer without re-configuring. Other user-friendly touches include a battery life icon that's always visible and on-line User Manuals.

6 FULL-SIZED KEYBOARD

Olivetti has exploited every available millimeter of width in giving you a keyboard that won't cramp your style. The letter and number keys are a full 17 mm on each side - like a

Comfortable keyboard

desktop. The built-in trackball is mounted at right beneath the display in an ergonomic, flexible base which coddles your thumb as you point the cursor. There are two clicking positions, one on the right rear corner and a second opposite the trackball for left-handers.

Ergonomic trackball

7 UNLIMITED FLEXIBILITY

Echos imposes no limits on your productivity. The HDU is field-service upgradable, and widely available standard SIMM RAM expansion modules of 4, 8 or 16 MB are a snap to install using only a screwdriver.

Full connectivity

The integrated PCMCIA Type III slot with Phoenix driver gives you universal connection availability: fax/modem, local network and host I/O, HDU enhancement, and more. Serial and parallel ports are standard, and a range of options includes external FDU, port replicator, VGA adapter and spare battery.

8 WORK AT 10,000 FT.

That's right. You can use your Echos up to a height of 3,000 meters. In addition to certifiying the reliability of its components, Olivetti subjects Echos to

Ultraslim design

a battery of tests to ensure that it meets the highest standards for resistance to shock, cold, heat, pressure, humidity, radio interference, vibration, noise and, of course, altitude.

9 TESTED WITH HUNDREDS OF HW/SW PRODUCTS

Olivetti's commitment to indu-

stry standards and compatibility is only fitting for a company that is the N° 1 supplier of multivendor service in Europe. Echos 20 is tested compatible with over 100 leading software packages and with the products of all major PCMCIA card suppliers. Use it today in the knowledge that you will be able to use it tomorrow, no matter how the I.T. market evolves.

10 CUSTOMER CARE

Olivetti's dedication to customer satisfaction can be seen in the two standard warranties applicable to every Echos, one national and one international. If your work takes you abroad, take your Echos. And your Echos will automatically bring its guarantee. For even more peace of mind, Olivetti Service Plus offers a wide range of warranty extensions and additional services to meet your individual needs. All this is supported by Olivetti's international network of 5,500 Dealers and Systems Partners and 15,000 technical specialists.

olivetti

How to Contact the Editors

We welcome your questions, comments, complaints, kudos, and submissions.
MAIN OFFICE: One Phoenix Mill Lane, Peterborough, NH 03458, (603) 924-9281.
San Mateo: 1900 O'Farrell St. #200, San Mateo, CA 94403, (415) 513-6912.
New York: 1221 Avenue of the Americas, New York, NY 10020, (212) 512-3175.
Lexington: 24 Hartwell Ave., Lexington, MA 02173, (617) 863-5100.
GERMANY/EUROPE: Liebigstrasse No. 19, 60323 Frankfurt, Germany, +49 69 7140 7123.
ELECTRONIC MAIL: On BIX, send to "editors." All BYTE editors and columnists also have individual mailboxes on BIX for easy access.
MCI: 250-0135 BYTE Magazine. Many editors also have individual MCI addresses in their own name.
OTHERS: Many editors also are reachable through uunet, AppleLink, CompuServe, and numerous other services.
U.S. fax: Editorial: (603) 924-2550
Advertising: (603) 924-7507
U.K. fax: +44 71 495 6734

SUBMISSIONS:
Authors: We welcome article proposals and submissions. Unacceptable manuscripts will be returned if accompanied by sufficient return postage. Not responsible for lost manuscripts or photos.
Vendors: We welcome news of your new products; please call the News department or the BYTE Lab at the earliest possible date. We cannot be responsible for unsolicited product samples.

ARTICLE REPRINTS:
For price quotations on customized reprints of BYTE articles, contact Susan Monkton, reprints manager, at (603) 924-2618. (Minimum quantity: 500.)

Subscription Customer Service

Inside U.S. (800) 232-BYTE; outside U.S. +609 426 7676. International subscribers may also contact our international customer service facility in Galway, Ireland, by calling +353 91 752792 or via fax: +353 91 752 793. **For a new subscription, (800) 257-9402** U.S. only, or write to BYTE Subscription Dept., P.O. Box 555, Hightstown, NJ 08520. Subscriptions are $29.95 for one year, $54.95 for two years, and $74.95 for three years in the U.S. and its possessions. In Canada and Mexico, $34.95 for one year, $64.95 for two years, $87.95 for three years. In Europe, £42 (US$60) for fast surface delivery, £55 (US$80) for air delivery. Non-European countries US$60 for surface mail, or US$85 for air mail. Single copy price is $3.95 in the U.S. and its possessions, $4.95 in Canada. Foreign subscriptions and sales should be remitted in U.S. funds drawn on a U.S. bank. Please allow six to eight weeks for delivery of first issue.

BIX Interactive On-line Service

BIX, owned and operated by Delphi Internet Services Corporation, is a worldwide, low-cost, on-line information service featuring industry news, downloadable software, powerful electronic mail, previews of upcoming BYTE articles, the full text of published issues of BYTE, and source and/or executable code for BYTE benchmarks and noncommercial software mentioned in feature articles. BIX also offers unmatched "conferences" on virtually every computer-related topic imaginable, where you can share information with thousands of other computer pros. To subscribe via modem, set your communications software to full duplex, 7 bits, even parity, 1 stop bit, and then call (800) 695-4882 or (617) 491-5410, or telnet to x25.bix.com and type "bix" at the USERNAME prompt. At the Name? prompt, type bix.ville. For more information, call (800) 695-4775 or (617) 354-4137 (voice); send a fax to (617) 491-6642; or send Internet mail to info@bix.com.

BULLETPROOF

Introducing the latest additions to the

We've just introduced a host of new ways to benefit from APC's award-winning reliability. With solutions from surge protectors through midrange UPS systems to protect anything from a basic PC to the most advanced superserver, APC is your one-stop shop for reliable power protection. Call today for a more information on our new, enhanced product line of "bulletproof" protection!

Use APC's new UPSs with our other award-winning products and accessories...

SurgeArrest® AC surge suppressors...

New ProtectNet™ data line surge suppressors...

Accessories including Measure-UPS™, Call-UPS™ II and Share-UPS™...

PowerManager™ power control centers...

NEW!

Back-UPS®

Back-UPS® Pro™

BEST VALUE POWER PROTECTION FOR PC WORKSTATIONS
Use with Compaq Presario, IBM ValuePoint and equivalent 386 and 486 PC workstations

PERFORMANCE PROTECTION FOR ADVANCED UNIX WORKSTATIONS
Use with SUN, HP, Silicon Graphics, IBM RISC 6000, and UNIX workstations

■ SIMPLE SIGNALING	■ SIMPLE SIGNALING

- ■ Unmatched surge and lightning protection for maximum hardware safety
- ■ Site diagnostics automatically spot missing ground and reversed polarity
- ■ LAN signaling allows simple shutdown with interface kits (400 and above)
- ► User replaceable, hot-swappable batteries insure uptime safe disposal
- ■ $25,000 lifetime Equipment Protection
- ■ Available in 200, 280, 400, 450, 600, 900, 1250 Volt-Amps. Starting at $119 list.

All the protection of Back-UPS, plus:

- ■ Automatic voltage regulation provides enhanced protection against long duration brownouts and overvoltages
- ■ Built-in network card/modem surge protection for maximum system integrity
- ► Intelligent Battery Management extends battery life and increases runtime
- ■ Line-interactive design for better performance and reliability
- ■ Available in 280, 420, 650 Volt-Amps. Starting at $199 list.

KEY:

Performance, price increase left to right

► **APC Exclusive**

Software capability rating (1 good-3 best)

■ Simple signaling provides safe shutdown in the event of an extended power loss. APC offers the industry's widest selection of kits, including PowerChute.

■ ■ *Serial signaling allows UPS queries, and data logging*

and analysis. PowerChute v/s offers serial signaling.

■ ■ ■ *Enhanced serial signaling and SNMP management allow full wide-area control, such as PowerChute PLUS and PowerNet.*

PROTECTION
world's best-selling UPS protection

NEW!

NEW!

Smart-UPS® v/s™

Smart-UPS®

Matrix-UPS™

BEST VALUE SERVER PROTECTION FOR SMALL BUSINESS NETWORKS

Use with PS/2, Dell, ProLinea and equivalent small business or department servers

SERIAL SIGNALING

All the protection of Back-UPS Pro, plus:

■ Serial communication for enhanced network shutdown and control. Allows power quality event log, scheduling and UPS self-test.

▶ Ships with PowerChute v/s software for maximum reliability and convenience

■ Available in 420, 650, 1000, 1400 Volt-Amps.Starting at $349 list.

PERFORMANCE POWER MANAGEMENT FOR ADV'NC'D BUSINESS NETWORKS

Use with Proliant, IBM PC 55, Alpha and equivalent servers and internet devices

ENHANCED SERIAL SIGNALING & SNMP MANAGEABLE

All the protection of Smart-UPS v/s, plus:

■ Software configurable features

▶ SmartSlots for SNMP manageability means wide-area reliability and uptime

■ Full sine wave output

■ Supports APC monitoring accessories, such as Share-UPS, Measure-UPS, Call-UPS and more.

■ Rackmount and extended run models

■ Available in 250*, 400*, 450, 700, 1000 1400, 2200, 3000 Volt-Amps. Starting at $299 list.* Not all features available on these units

ADVANCED MODULAR PROTECTION FOR DATACENTERS

Use with AS/400, HP9000, VAX, NetFrame and equivalent minis-servers

All the protection of Smart-UPS, plus:

▶ Modular hot-swappable design for 100% uptime

■ Add runtime quickly and easily with microprocessor-controlled SmartCells

■ 60% lower lifecycle costs than comparable brands.

■ Available in 3000 and 5000 Volt-Amps. Starting at $3499 list.

Solutions to suit any application

APPLICATION	VALUE	PERFORMANCE	HI-PERFORMANCE
WORKSTATIONS	BACK-UPS	BACK-UPS PRO	SMART-UPS
SERVERS	SMART-UPS v/s	SMART-UPS	MATRIX-UPS
DATACENTERS	--	SMART-UPS	MATRIX-UPS

Circle 62 on Inquiry Card.

APC™

AMERICAN POWER CONVERSION
Dept. A2

800-800-4APC

Responding to Rafe

I'm taking Raphael Needleman up on his invitation in "Information Underload" (February Editorial). I strongly suggest that you give BYTE a significant presence on the WWW (World Wide Web). Several other magazines are doing so; BYTE could do it better.

Jim Tubman
Research Officer
Calgary, Alberta, Canada
jbtubman@skyler.arc.ab.ca

As a wildly enthusiastic computer user and aggressive info-hound, I have to say you've written an editorial with a number of good points. However, one man's garbage is another man's food. We all need filters the size of a baleen. There would have to be 50 of me to absorb it all at the depth I'd like to. We need many clones or multiple lifetimes.

Lee van Laer
71623.277@compuserve.com

Chintzy on the Test Results

If the intent of the January Lab Report was to answer the question, "Which is the best monitor among the 17-, 20-, and 21-inch classes?", it missed the boat. Many have "excellent" scores in the best overall, quality index, and features columns. How can we differentiate monitors with the same scores? Why not publish the test results and let the reader decide? We want to know which monitor has the best picture—period. If you answer this question at all, you do it without elegance, clarity, or quantitative precision.

Vic Spelman
victor@alleys.com

An Unfocused Rant

In "Networks for the Enterprise" (February), the author rates Windows NT above OS/2 LAN Server. The Workstation version of Windows NT 3.5 is too slow to use on a 16-MB 486. I had to increase my swap file size from 15 to 32 MB to make Microsoft Visual C++ 2.0's resource compiler compile a small resource file on the command line. The operating system, the SDK (Software Development Kit), and Visual C++ 2.0 together take nearly 400 MB of disk space. The minimum 4-MB RAM requirement should be 40 MB.

Problems with Windows for Workgroups 3.11 include "internal error" messages and corrupted NetWare net.cfg files. Microsoft orphaned Win32s in C++ 2.0; Win32s APIs are no longer documented. You cannot compile or debug Win32s applications under Windows for Workgroups 3.11. The only debugging tool remaining crashes a lot, and Windows NT 3.5 breaks Visual C++ 1.1's 1.0 linker. When Microsoft helped against IBM's monopoly, we all sided with Microsoft. Now Microsoft is the monopoly, and it shows.

Xiaodong Tan
Xtan@prairienet.org

We propose many reasons why OS/2 LAN Server 4.0 did not perform as well as Windows NT Server 3.5: Its client software doesn't support the protected-mode redirector of Windows for Workgroups 3.11; we had to test LAN Server in 16-bit mode file I/O; and the client software would not run with Windows for Workgroups' 32-bit disk and file I/O enabled.

Another factor could be that LAN Server 4.0 does not provide a configurable volume block size, which should match the hardware's blocking factor for disk striping. We attached a Core International disk array that has a default configuration of 32 KB. Unlike other LAN operating systems that provide volume block size options (e.g., you can set NetWare's block size from 4 to 64 KB), we did not know LAN Server's block size and could not change it anyway.
—Tadesse Giorgis

Apple's Newton MessagePad Phones Home

I find BYTE's coverage of late-breaking technology and its practical applications excellent reading. Andy Reinhardt's article "The Newton Goes Vertical" (January) on the customization of Apple's Newton MessagePad into a field service machine for telephone technicians is extremely timely. Several weeks after the article came out, I mentioned it to a telephone technician who was using such a machine at my workplace. He gave me a first-class tour of the device, and it appears to live up to Reinhardt's technical specifications.

G. Morrison
Address withheld

I just want to compliment Andy Reinhardt on the Dracon CDA (craft digital assistant) piece in the January issue—virtuoso reporting and writing! I have never been so mesmerized when reading about a test widget. In addition, my mother could have understood it. Well done.

Paul McCloskey
paulmcc@mcs.net

OS/2 Drivers for Graphics

Why did February's Lab Report, "True-Color Graphics Accelerators," fail to cover any OS/2 driver support or benchmark results?

Peter Skye
Glendale, CA

OS/2 testing was a part of our original test methodology. But when we began testing last September, many of the adapters were brand-new, and few arrived with OS/2 drivers. So we requested them, but of those we received, most were not ready for prime time. In Barry Nance's review of Warp (March, page 131), you'll find a table of graphics adapters and chip sets that Warp supports.—Eds.

Benchmarks

The BYTE/NSTL Lab Report "True-Color Graphics Accelerators" (February) was great. It was nice to see a measure of video-playback performance. Because I'm working on a graphics/video subsystem under Windows 3.1/95 for Digital Equipment, I'm interested in your InterMark video and graphics tests used throughout the article. Are the InterMark tests externally available?

Andrew Wilson
DCI Developer
wilson@pasta.enet.dec.com

If you think Exabyte is just 8mm, think again. Today, Exabyte is Everywhere.

- ► High capacity
- ► Superior reliability
- ► Software compatibility
- ► Worldwide support
- ► Tape media products
- ► Certified service

EXABYTE IS EVERYWHERE

8 MILLIMETER

4 MILLIMETER

QUARTER-INCH

LIBRARIES

It's easy to think of Exabyte as the company that pioneered helical-scan 8mm tape backup and restore technologies. But today, Exabyte is much more. We offer customers worldwide a full line of backup and restore products—8mm, 4mm and quarter-inch minicartridge tape drives, automated tape libraries, certified media products and service—all with the Exabyte name.

We're not just 8mm anymore. Today, Exabyte is Everywhere.

EXABYTE®

In North America, call our office at 1-800-EXABYTE. In Europe, call The Netherlands office at 31-3403-51347. Or in Asia, call our Singapore office at 65-2716331.

©1995 Exabyte Corporation. Exabyte is a registered trademark and "Exabyte is Everywhere" is a trademark of Exabyte Corporation. Exabyte Corporation, 1685 38th Street, Boulder, Colorado 80301 USA. Phone (303) 442-4333.

Circle 73 on Inquiry Card (RESELLERS: 74).

RFA (Request for Algorithm)

I found John Cuadrado's "Mining Statistics" (February) very interesting. My students and I are working on a problem for the Navy related to the collection of medical symptoms data from military personnel. The Navy is aware of algorithms and expert systems that take this data and output a probable diagnosis. What they want is a method of aggregating the data sequentially (as it comes in on a patient-by-patient basis) into a few continuously monitored measurements. They want to identify increasing patterns of illness quickly.

My idea is to use the geographical information on endemic diseases to establish known centroids in "symptom" space. The hard part is that a significant number of patients that cannot be classified into one of the endemic diseases might accumulate. We need an algorithm that detects when a new cluster is required and then adjusts the classification scheme accordingly.

John Angus
Professor of Mathematics
The Claremont Graduate School
Claremont, CA
angusj@cgs.edu or angusjohne@aol.com

Solutions-R-Us

Having read Rick Grehan's "Seer-Sem Offers Realistic Forecasting for Programmers" (September 1994, page 34), which describes Galorath Associates' expert system for estimating the size, cost, and resources of a development project, we are now negotiating with Galorath for distributorship of Seer-Sem in India. Thank you for the tip.

A. Dear
The Idea Machine
P.O. Box 4268, India
dear.idea@axcess.net.in

An Incomplete Definition

On page 72 of the February issue, Oliver Sharp defines numerical analysis as the study of errors caused by computer arithmetic. Not so! The job of numerical analysts is to develop and analyze algorithms for mathematical problems in science and engineering. A numerical algorithm must produce an answer to a problem accurately and efficiently.

Consider two recent breakthroughs in the field: Karmarkar's interior-point method for linear programming and Greengard and Rokhlin's "fast multipole" algorithm for n-body problems. In each of these, unexpected insight led to a new algorithm significantly faster than its predecessors. Analysis of round-off error played no role in these discoveries.

Stephen Vavasis
Department of Computer Science
Cornell University, Ithaca, NY
vavasis@cs.cornell.edu

Programming in Prolog

Dick Pountain's "Constraint Logic Programming" (February) propagates the myth of Prolog's poor execution performance as the main hindrance to its commercial acceptance. The best Prolog systems compete in performance to C (e.g., in some cases, the Aquarius system, distributed by USC, executes faster than C). Pountain identifies Prolog's search process as the root of its inefficiency, but you can easily program efficient search tools with Prolog. The key is to use the built-in search only as a primitive when developing new search algorithms.

Peter Van Roy, Ph.D.
Saarbruecken, Germany
vanroy@dfki.uni-sb.de

Silver-Streak Sauté

The following oversights occur in Stan Miastkowski's "Daisy-Chain Ethernet" (January): Silver-satin telephone line cord is 26-gauge stranded, not 18 gauge. The lower the wire gauge number, the thicker it is. Copper wire has a series of frequencies that are multiples of the pulse rate, not memory. When an NIC's (network interface card) impedance is not the same as the cable's, the different frequencies travel back and forth at different speeds and attenuation levels. This causes problems such as ringing and overshoot. Impedance is not specified on telephone wire, so you cannot design an NIC to match it. However, you can use Schottkey diodes to clamp overshoot, the way surge supressors clamp spikes on power lines.

Finally, the nominal ringing generator voltage in analog telephone systems is 90 VAC 20 Hz. On-hook battery is nominally 48 VDC dropping to around 10 VDC off-hook. Where did the notion of 180 VDC come from?

Wilton Helm
Address unknown

Your comment on wire gauge is correct; a Silver Streak designer told me copper wire has memory, but your technical explanation is more accurate. Your comments on ringing voltage for standard analog telephone lines are also correct; what I presented describes a worst-case scenario for proprietary in-house systems.—Stan Miastkowski

FIXES

In "The Great Little File System" (February), we omitted Veritas' contact information. Veritas Software ((415) 335-8000) is in Mountain View, California.

The "Best of Comdex" text box (February, page 24) continues on page 28, not on page 30 as published. ∎

COMING UP IN MAY

- **QUICKTIME VIRTUAL REALITY**
  A new dimension of full-spectrum photo-realistic graphics from Apple and Microsoft.

- **DIGITAL VIDEO**
  Networking, compression, and how to set up your own video studio.

- **SUPER SERVERS**
  Mainframe capability in small packages.

- **PRINTERS**
  29 printers that mean business.

- **THE BRITISH LIBRARY IS ON-LINE**
  You can now search and retrieve text from OPAC (Online Public Access Catalog), the British Library's catalog of its 13 million volumes.

Choose Your Weapon.

Plextor's 4PleX Quad Speed CD-ROM Drives. Now Available in 256K and 1MB Buffer.

It's a war out there. And if you're engaged in the battle for multi-media performance or fast data access, your first line of attack is a quad speed CD-ROM drive from Plextor. Designed to plow through reference data, archives, and today's most sophisticated multimedia applications, the 4PleX line of reliable drives boasts an arsenal of features including a 600KB/sec transfer rate, access to the world's largest CD-ROM drive buffer, blinding access speeds, the security of a two-year warranty and unlimited toll-free technical support.

Plextor offers you a choice of weaponry. The 4PleX 256 provides quad speed multimedia performance with fast access to key data. To triumph over new high-demand video applications and games, or to impress your allies with sophisticated multimedia presentations, choose the original 4PleX

with a 1MB buffer. Like a reconnaissance scout, the buffer looks ahead and anticipates your needs. As a result, every frame of video advances through your computer. Both drives give you smooth, fast, winning performance without over-taxing your computer in strategic multi-tasking situations.

No wonder *PC Computing* gave the 4PleX a four star ranking and the title "King Quad." Today's CD-ROM explosion promises to take no prisoners. So arm yourself with a 4PleX and rest easy knowing you are equipped to win the battle. For your nearest Dealer and a free copy of our latest book "18 Questions to Ask Before Purchasing a CD-ROM Drive" call 1-800-4PLEXTOR (1-800-475-3986).

PLEXTOR

Circle 126 on Inquiry Card (RESELLERS: 127).

Mac Clones —Finally

Now, more than 10 years after Apple introduced the Mac, the first Mac OS–compatible systems sanctioned by Apple should be available soon

DAVE ANDREWS

By the end of this month, at least three companies expect to be selling Mac OS–compatible systems. Each company will use a different strategy to complement Apple's line of Macs. For example, the first generation of systems from Radius (Sunnyvale, CA) will consist of high-end Macs for video-editing and color-publishing professionals. And other companies will introduce highly expandable systems with extra NuBus expansion slots or low-cost systems built using PC manufacturing facilities and some PC components. Also, companies such as Pioneer will release Mac clones for the entertainment market.

Bridgette (El Cajon, CA), a vendor of mass-storage products for the Mac, will be one of the first vendors to sell Mac clones. The company's first Quatro 850 systems, which will be sold under the Cutting Edge label, will use Apple's Centris 650 motherboard, the Apple SuperDrive floppy drive, and either a 25-MHz 68LC040 with no FPU or a 33-MHz 68040.

Bridgette hopes to introduce systems lacking from Apple's product line by introducing midrange tower Macs with three NuBus slots (the only systems available from Apple with three NuBus slots are the Power Mac 7100 and 8100), one PDS slot,

Radius's Mac OS–compatible systems will be high-end Macs for graphics or video-editing professionals. Radius's first systems will use the same motherboard as Apple's Power Mac 8100/110 and a PowerPC 601 CPU.

five drive bays, and up to 128 MB of RAM. Future systems from the company will offer seven drive bays for RAID configurations.

The company may also offer DOS/Windows-compatible configurations through adapter cards from Orange Micro (Irvine, CA). For example, it ex-

pects to sell a 68LC040 version of the Quatro 850 with 8 MB of RAM, a 540-MB SCSI hard drive, a double-speed CD-ROM drive, and a 17-inch monitor for $1999. Another system, with 8 MB of RAM, a 270-MB SCSI hard drive, a keyboard, a mouse, and software utilities (but no monitor), will sell for under $1000.

Power Computing, a startup company from Milpitas, California, is using the same components as Apple but in a motherboard of its own design. It will sell its own brand of Mac clones as well as motherboards to companies that want to build Mac clones. Power Computing declined to release specific details of its systems, but sources say the company will use widely available PC components (e.g., power supplies and chassis) whenever possible to make less expensive Mac systems. Power Computing's systems should sell for 10 percent to 15 percent less than a comparable Mac.

Power Computing's first PowerPC 601–based systems will likely ship this month in two configurations, a tower and a baby AT chassis, with standard Mac connectors, such as an ADB (Apple Desktop Bus) port for the keyboard and a 15-pin video connector. The first systems will contain NuBus slots, with support for PCI (Peripheral Component Interconnect) to follow. The company is employing a number of tactics to keep its systems inexpensive, sources said.

By making its motherboard fit the baby IBM AT form factor, Power Computing makes it possible for PC vendors, some of which may have excessive manufacturing capacity, to easily build Mac clones. PC makers can simply drop Power Computing motherboards into the baby AT chassis of their choosing.

CompuAdd (Austin, TX), a vendor of desktop PCs, notebooks, and servers, has already

said that it will build systems for Power Computing that will be sold with the Power Computing label. CompuAdd expects to build 2000 to 3000 Mac clones a month.

Sources also indicate that Power Computing will reduce cost by building its motherboards using a four-layer process. In general, manufacturers make motherboards with four, six, or eight layers. Apple's Power Mac motherboards have six layers. Reducing the motherboard footprint generally requires more layers. Smaller footprints push the motherboard's components closer together, which means more copper wires are closer together. Adding layers insulates those wires from interfering with each other. But increasing the layer count adds manufacturing steps and translates to increased cost, says Jim Turley, who is senior editor of the *Microprocessor Report* (Sebastopol, CA).

Sources say that by using a four-layer process, Power Computing can reduce its cost by $20 to $25 per motherboard compared to Apple's motherboards. "Using the AT chassis is a really shrewd idea," Turley says.

Tests done on prototypes of Power Computing's desktop system using the 680x0 version of BYTE's benchmarks indicate an overall performance comparable to that of Apple's Power Mac 7100/80 for Power Computing's 80-MHz desktop. BYTE's tests also indicate that Power Computing's 100-MHz tower system turned in a per-

formance identical to that of a Power Mac 8100/100. BYTE will review the final systems in a later issue.

Volume purchasers of Macs look forward to Mac clones. "Pricing is important," says Theresa Woodward, who is a purchaser for Dartmouth College (Hanover, NH), which currently has over 8000 Macs and buys 1500 more annually for new students. "If the systems are 100 percent compatible [and cost less than Apple's], we're more than willing to go with the clones."

Developers are waiting to see how well Apple's strategy succeeds before committing new programming resources. "The Mac is currently a lower priority [in Lotus Development's presentation-graphics products] than Windows and OS/2" says Nina McIntyre, general manager of Lotus's graphics products group. "It's definitely a wait-and-see situation."

SERVER TRENDS

Low-End Servers Grow Up

It's easy to tell the difference between low- and high-end Intel-based LAN server systems: The processor, installable RAM, cache size, and hard disk capacity are usually bigger or better in the high-end servers. However, other vital features are becoming more affordable for vendors to implement across their entire server lines as more users demand low-end (under $5000) servers. As a result, high-end features such as fault tolerance, disaster recovery, and remote administration are migrating to the low end.

"Features that were typically found in high-end servers used for mission-critical applications are being demanded in less expensive servers," says Mary McDowell, director of systems product marketing for Compaq Computer (Houston, TX). Compaq stirred up the low-end market with its mid-February announcement of the ProSignia 300, a system that will start at $3200 and will include a 75- or 90-MHz Pentium configuration, ECC (error-correction code) memory, Compaq's TriFlex/PCI System Architecture, SNMP-based server management software, and other features typically associated with the high end.

PC-server vendors predict these features will migrate from the high end to the low end:

SMP (symmetric multiprocessing). Compaq recently replaced the ProLiant 1000, a single-processor model, with the SMP-upgradable ProLiant 1500. Vendors' future models

Server Sales Go Up

Units (millions) / **Revenue** (billions of dollars)

Source: Dataquest (San Jose, CA)

Actual and predicted worldwide server (all types) sales, in terms of units and revenue. For 1994, figures are estimated. For 1995 on, they're forecasted.

are bound to bring the SMP option down to even less expensive models.

Cache type. Closely aligned with SMP is the cache issue, because servers that acquire multiprocessing capability will initially do so using a shared cache. Some examples are the ProLiant 1500 and AST Research's Manhattan P series. Not expected to filter down this year is a dedicated cache per CPU, but there's no slowing the growth of the shared cache.

ECC memory. ECC is now becoming a checklist item for low-end server buyers. IBM, Hewlett-Packard, and others will probably follow Compaq's lead in providing ECC across the board, offering it at least as an option. Marketing wars may ensue when schemes like Compaq's Advanced ECC—which corrects up to 4-bit errors—migrate, perhaps this year, to the low end.

Disk arrays. For small LAN servers with low CPU utiliza-

tion, a software-based disk array uses the server CPU to stripe redundant copies of data across arrayed drives. If one disk fails in a RAID-5 array, the other drives can maintain the data. But hard drives and the more efficient but expensive hardware controllers, which use a RISC processor for the striping, are also coming down in price.

Manageability. IBM's NetFinity, Compaq's Insight Manager, and HP's NetServer Assistant compete at the high end, and wherever possible, vendors are bringing features such as remote management and replicated software installation to the low end. As workgroup servers add application serving to their current file and print service responsibilities, companies need to remotely manage and monitor servers at remote locations that don't have on-site database administration personnel.

—Ed Perratore

GRAPHICS MEMORY ARCHITECTURES

Faster Graphics Cards on the Horizon

PC buyers will see a tremendous leap in graphics performance at the end of this year or early next year, thanks to new kinds of graphics memory. In the past, graphics chip makers increased performance by doubling the number of bits that the graphics controller could process at once (e.g., moving from 32- to 64-bit-wide graphics architectures). Although Number Nine Visual Technology's (Lexington, MA, (617) 674-0009) Imagine-128 ($999) graphics card uses a 128-bit graphics controller, other card vendors believe that doubling the controller interface past 64 bits is no longer cost-effective.

Dean McCarron, an analyst with Mercury Research (Scottsdale, AZ), a PC chip marketing and consulting firm, says the number of pins required and the size of the die are so large, among other problems, that the price/performance ratio of 128-bit controllers is not compelling.

Vendors are using or evaluating new memory architectures that offer better performance than DRAM but cost less than the usually more expensive video memory. Analysts predict that the change will improve graphics performance by 400 percent or more over current DRAM solutions in low-end and midrange PCs, with a markup of about $9 for low-end card manufacturers. Users should see prices for video-memory-type performance drop slightly while performance increases, improving motion video quality and making it possible to integrate 3-D acceleration into regular graphics controller chips.

It's too early to tell which of the many new memory architectures—there are at least eight—will gain enough adherents to lower prices significantly. PC and card vendors must evaluate other factors than performance alone. Also in the mix are cost and availability of chips and second sources. EDO (Extended Data Out) DRAM and SDRAM (synchronous DRAM) have enough backers to ensure that they will virtually replace standard DRAM on low-end and midrange PCs this year, manufacturers and analysts say.

EDO DRAM chips, which are available from several manufacturers (e.g., Micron, NEC, and Samsung), offer about 20 percent faster performance than that of standard DRAM. This type of memory is already being used interchangeably with standard DRAM by board manufacturers. By the end of the year, it will be widely used and as inexpensive as standard DRAM, says Jim Handy, director and principal analyst for Dataquest (San Jose, CA). EDO DRAM gets its performance improvement by loading data at the same time it is switching to a new address, rather than waiting between these operations, as standard DRAM does.

SDRAM approaches video-memory performance. It can send data faster because all its parts are synchronized, or tuned to a clock, so the system doesn't have to wait as long between data accesses to ensure the integrity of the information, as standard DRAM does. Many graphics manufacturers have endorsed SDRAM. Prices will drop enough that it will add no more than a 10 percent to 20 percent premium over standard DRAM, McCarron says.

Whether speedier, more exotic forms of memory such as MDRAM (multibank DRAM) and RDRAM (Rambus DRAM) will achieve high enough sales to let their prices compete with standard DRAM or video memory is still undetermined. MDRAM and RDRAM are both capable of achieving speeds faster than those of video memory. Both memories are single-ported, allowing their full bandwidth to be allocated dynamically where it's needed, rather than always keeping half in reserve to refresh the screen, as is the case with video memory.

MoSys (San Jose, CA), the company that makes MDRAM, and Rambus (Mountain View, CA) have promised their memories will sell for no more than a 20 percent premium over standard DRAM. Whether this is achievable depends on how much demand there is for their products, McCarron says.

While many companies have praised MDRAM for its performance, they have not rushed to adopt it. This is because MoSys is new and hasn't yet announced a second source, says Jon Peddie, principal of Jon Peddie Associates (Tiburon, CA), a consulting company that specializes in computer graphics hardware. Rambus has several suppliers. RDRAM has been endorsed by Cirrus Logic, a supplier of graphics memory controllers. It will also be used by Nintendo in its forthcoming Ultra 64, a 64-bit video-game system. The Nintendo relationship should boost Rambus's volume sales.

A new type of dual-ported memory from Samsung, known as WRAM (Window RAM), will offer performance that's up to 50 percent better than that of the fastest video memory available today, Matrox officials claim, but at a price that's 20 percent lower per bit. Analysts don't expect WRAM to drop to DRAM prices because its dual-ported architecture requires a larger and more expensive die.

—Cate T. Corcoran

MEMORY TECHNOLOGIES COMPARED

The chart shows projections of price for late this year and early next year, along with the peak bandwidth supported by these architectures. The chart also shows which companies had declared support for these memory architectures at press time. Some of the bandwidth numbers are theoretical. Many vendors currently use 64-bit DRAM. (Sources: prices, Mercury Research; peak bandwidth, Jon Peddie Associates.)

MEMORY BUS WIDTH	PEAK BANDWIDTH (MBPS)	COST FOR 2 MB	WHO'S USING IT
64-bit DRAM	200	$45	Chips & Technologies
64-bit EDO DRAM	320	$47	Chips & Technologies, ATI Technologies, S3
64-bit SDRAM	530	$49–$53	Chips & Technologies, Number Nine, Tseng Labs
Two 8-bit channels of RDRAM	1000	$49–$53	Cirrus Logic, Silicon Graphics
64-bit MDRAM	800	$49–$53	Trident, Tseng Labs
64-bit WRAM	800	$68	Matrox Graphics, Trident
64-bit VRAM	600	$65	S3, ATI Technologies

NETWORKS

Microsoft Furthers NetWare-to-NT Transition

Microsoft continues to ease the transition for users migrating from NetWare to Windows NT Server with its File and Print Service for NetWare, or FPSNW. Currently undergoing beta testing, with a shipment date dependent on customer feedback, the FPSNW utility makes Windows NT Server look like a NetWare 3.x file and print server.

FPSNW should appeal to companies that are tapping Windows NT Server as an applications server to host mission-critical programs such as SQL Server but don't want to maintain another NOS (network operating system) server—NetWare—for file and print services. With FPSNW, companies can deploy one server to handle file, print, and applications services without having to change their current NetWare client software.

FPSNW is the latest in a series of Microsoft initiatives to offer a migration path from NetWare or simply use NT Server on NetWare LANs, explains Jamie Lewis, president of the Burton Group (Salt Lake City, UT), an information services firm that specializes in network computing. Prior to Windows NT, Lewis explains, Microsoft required you to run SQL Server on its OS/2-based LAN Manager. That strategy backfired because most users weren't willing to ditch their NetWare investment just to run SQL Server. "Microsoft realized they had to accommodate reality," Lewis says.

The first step in those accommodations was Microsoft's addition of support for NetWare's IPX/SPX protocol to Windows for Workgroups and Windows NT Server, which allowed for such configurations as a computer running Windows NT Server and SQL Server on a NetWare LAN. Microsoft also added NetWare client support to Windows NT Workstation and the forthcoming Windows 95, which let those operating systems use the file and print services of a NetWare server.

FPSNW, says Lewis, "is for people who want to get away from NetWare servers, period."

With Microsoft's File and Print Service for NetWare, users can log on to a Windows NT server that looks just like a NetWare server. The utility lets companies consolidate on Windows NT and upgrade their NetWare clients at a later date.

Prior to FPSNW, eliminating NetWare servers would require LAN administrators to throw away their investment in NetWare clients. With FPSNW, when users log on to a Windows NT server, it looks the same as a NetWare server. This way, companies that want to eliminate NetWare servers can do so right away and replace the client software when they are ready.

FPSNW alone probably won't convince users to choose NT Server over NetWare. But it's easier for companies with NetWare installations to choose NT as an applications server now and replace the NetWare client later (e.g., when Microsoft releases Windows 95).

—Jon Udell

INTERNET

Free Rides Are Disappearing

Unisys recently decided to collect license fees for the data-compression algorithm used in GIF, a popular graphics standard of the on-line world. This illustrates how certain Internet components that users now take for granted may eventually cost them money.

Seven years after its patent on the LZW (Lempel-Ziv-Welch) data-compression algorithm had been granted, Unisys (Blue Bell, PA) began requiring licensing fees for programs using LZW. CompuServe's (Columbus, OH) GIF uses LZW. Last year, developers in the graphics and on-line communities expressed dismay when CompuServe announced that as a result of Unisys's request for licensing fees, CompuServe would start charging fees for programs that use GIF images.

For years, GIF, while copyrighted by CompuServe, had been a de facto standard as well as a popular graphics format for graphics and desktop publishing. Faced with a storm of disapproval from developers and users, CompuServe said it would coordinate the development of a non-LZW follow-up format called GIF24, which will be a 24-bit, lossless open graphics format.

This controversy made users more wary of informal open standards and helped dispel the notion that the Internet and its resources are free. Jeremy Allaire, who is president of New World Media (Minneapolis, MN), a company that tracks trends and events in on-line publishing, says that users who assume that everything on the Internet is free are "sadly mistaken."

Some costs associated with the Internet, such as network setup, can be hidden from users when their employer pays them. Other costs of using the Internet will become evident to end users due to supply-and-demand issues. Users who can't afford to wait the half hour or so it can take to access a Veronica or Web searching site during its busiest hours may have to turn to surcharged Internet services.

Meanwhile, developers are exploring GIF replacements, including Graphics Exchange Format (a version of GIF that doesn't use LZW), the Portable Bit-Map Format, and the Free Graphics Format. Other developers are exploring the use of preexisting graphics formats in their applications, such as the open JPEG format.

—Steven J. Vaughan-Nichols

INTERNET

The Net's Next Big Thing: Virtual Reality

A platform-independent standard for VR (virtual reality) called VRML (Virtual Reality Markup Language) could make navigating through on-line museums, libraries, and marketplaces on the Internet as common as interacting with textual information is on the WWW (World Wide Web) today. By the time you read this, VRML pioneer Mark Pesce expects to have the first freeware VRML authoring tools and browsers for Windows, Mac, and Motif platforms available on the Internet (http://www.eit.com/vrml/).

Most of the current efforts—including VRML—revolve around nonimmersive (i.e., special headsets or helmets are not supported) VR.

VRML 1.0 is based on a subset of Silicon Graphics' (Mountain View, CA) object-oriented OpenInventor 2.0 3-D scene and graphics description file format. Developers use VRML 1.0 to create objects that are rendered as an Internet user "walks" through a virtual room. Developers can attach hot links to HTML (Hypertext Markup Language) so users can click on an object and view additional text or images.

VR on the Internet does not require VRML. Using in-line graphics—good old-fashioned Mosaic with pictures in it—you can view graphics in the middle of a Mosaic page. By pushing a button that moves the point of view forward, backward, or sideways, the graphics are updated as you "move" through space. Each selection causes a network request to the Mosaic server that renders a new scene and/or loads a new HTML page.

However, Sandy Ressler, who is project leader of the Open Virtual Reality Testbed at the National Institute of Standards and Technology (Gaithersburg, MD), says the in-line method's drawback is performance. "You don't get the real-time response of VRML," he says. VRML is file-based, and once a 3-D scene is transferred to your machine, performance is as fast as your local PC can deliver. Also, because VRML is based on OpenInvention, which sits above Silicon Graphics' OpenGL graphics library, OpenGL

VRML browsers from companies such as Template Graphics Software (San Diego, CA) will let Internet users interact with nonimmersive VR applications.

accelerator cards will accelerate VRML.

The prospect of widely available VRML browsers received a big boost when Silicon Graphics published the source code for parsers of VRML so developers can write VRML viewers. Companies such as Netscape Communications (Mountain View, CA, (415) 528-2619) and Template Graphics Software (San Diego, CA, (619) 457-5359, http://www.sd.tgs.com/~template) plan to integrate VRML into viewers.

Companies such as TGS are also developing VRML authoring packages. Rik Carey, manager of 3-D graphics software at Silicon Graphics, notes that writing tools that convert a file from a specific modeling application's file format to the VRML format should be painless. "If you can write an authoring tool, you can convert it to VRML," he says. Other toolkit developers are taking a wait-and-see attitude.

"My feeling is a lot of very good ideas have gone into VRML, and it may be a standard that follows HTML into the big league," says Kevin Yurica, product manager for Autodesk's (San Rafael, CA, (415) 332-2344) Cyberspace Developer Kit, which is a C++ class library for developing VR applications and simulations for Windows and Windows NT. "CDK is consistent with the direction that VRML is moving toward, but there are some issues that have to be addressed."

One of those issues is the lack of support in VRML 1.0 for assigning interactivity attributes to 3-D objects. For example, with VRML 1.0, you can walk around and look at a 3-D representation of three balls atop a table, but you can't currently push the table over and watch the balls crash into each other. Pesce and others say a future version of VRML will add that support, along with support for letting multiple people interact with each other in a virtual world.

—John Vacca

NETWORKED VR FROM A CAVE

Most Internet VR (virtual reality) projects don't support immersion, where you wear a helmet and gloves and are immersed in a simulated environment. But researchers at the Electronic Visualization Laboratory at the University of Illinois at Chicago ((312) 996-3002) have developed a new model for networked VR, called a Cave.

A Cave is a 3-D environment that consists of a 10- by 10- by 9-foot room with three rear-projection screens for the walls and a down-projection screen for the floor. Instead of wearing a helmet, viewers wear LCD stereo shutter glasses that separate the alternate fields going to the eyes. Images on the screen "move" with the viewers and surround them. An image's direction and positioning are determined by a person wearing a magnetic tracking device from Polhemus (Colchester, VT). The computers supplying the power behind a Cave are from the Silicon Graphics Supercomputing Systems Division. They can handle the billions of FLOPS required when interacting with complex representations of weather patterns, seismic activity, or industrial CAD/CAM designs.

Tom DeFanti (Tom@uic.edu), the director of EVL, says that so far, researchers exploring networked Cave-to-Cave simulations have proven the concept of remote person-to-person immersive reality by performing basic actions (e.g., one person "touches" another person's head) over the Internet using a T1 line. For more complex applications, DeFanti is waiting for the activation later this spring of a VBNS (Very high Bandwidth Network Service) line that will dedicate a 155-Mbps ATM (asynchronous transfer mode) line between two supercomputer centers.

A Cave represents the high end in networked VR. DeFanti notes that solutions like VRML can't address the needs of these high-end VR applications. "If you need to sit in a simulation of a car, you want to reach out and see that everything is in the right place," he says. "You can't do that in a workstation." As you might expect, the price of a Cave is high. But DeFanti says that for some companies, the price is worth it. "It's beneficial enough that several of our clients are building Caves," he says. "Considering that the price is about $1 million per Cave, there's obviously some real communication going on."

—John Vacca and Dave Andrews

So fast, it will push your eyeballs to the back of your head.

The Tadpole P1000 is the *fastest* notebook on the planet. Hands down. In fact, with **100MHz** of pure Pentium™ power and PCI bus, it's probably faster than any workstation that's in your office right now. Add up to 810MB disk capacity and 128MB of RAM to the P1000 for stellar performance any time you need it.

The P1000 was built for sheer speed, but it was designed to handle the ROAD comfortably too. Its Lexmark KEYBOARD with integral pointer and three mouse buttons is a nice touch. With 1MB of video RAM, large 10.4" active-matrix TFT TRUE COLOR

screen and 16-bit CD-quality sound, every presentation is a BIG PRODUCTION. And to make a case for durability, we made a lightweight, rock-hard magnesium shell to protect the P1000 from life's hard knocks. Best of all, the processor module is factory-upgradable, and the SCSI-2 removable hard drive and memory are user-upgradable. This is the LAST notebook you'll ever need to buy.

Although the side effects of using the P1000 may be slightly disconcerting at first, just think of the corporate advantages of having your eyes in the back of your head. Call **1-800-232-6656** today!

intel inside

pentium PROCESSOR

NSTL
MEETS TEST STANDARDS
DOS/WINDOWS COMPATIBLE

TADPOLE TECHNOLOGY

United States: 512 219 2200
Fax: 512 219 2222
United Kingdom: 01223 428200
Fax: 01223 428201
France: 1 60 86 27 92
Fax: 1 60 86 27 89

Intel Inside and Pentium Processor Logos are trademarks of the Intel Corporation. All other trademarks used herein are property of their respective corporations. ©1995 Tadpole Technology.

Circle 117 on Inquiry Card (RESELLERS: 118).

WINDOWS SUITES

PerfectOffice a Strong Contender

WordPerfect's PerfectOffice is a latecomer to the Windows applications suite market, following Microsoft Office and Lotus SmartSuite. But an informal survey of PerfectOffice users indicates that WordPerfect (Orem, UT, (801) 225-5000), Novell's Applications Group, has a strong contender. This is especially so for users who want to integrate several applications in doing a single task.

PerfectOffice includes the WordPerfect 6.1 word processor, the Quattro Pro 6.0 spreadsheet, the Presentations 3.0 presentation-graphics program, the InfoCentral PIM (personal information manager), the Envoy portable document viewing and publishing solution, and a GroupWise client, for E-mail, calendaring, and scheduling clients. A professional version adds the Paradox windows database and AppWare, a visual development tool. We interviewed seven PerfectOffice users to determine their favorite integration features:

PerfectScript: This task-automation language, which lets users record macro commands across applications without having to learn a programming language such as Visual Basic, was a big hit with users. PerfectOffice, which is similar in its syntax to Microsoft's Visual Basic, also lets you write macro commands from scratch. But several users, such as Bruce Norton, owner of Norton Innovation (Lititz, PA), a VAR and consultancy, said PerfectScript's strength is its record feature that lets you quickly start automating tasks across applications.

Envoy: Users were intrigued by the possibilities of Envoy. It lets you print a document to a portable file format and distribute the file to Mac or Windows PCs users so they can view and annotate the document. "I think Envoy has promise in the area of workgroup collaboration," says Paul Dalton, partner at the law firm Jackson & Walker, which has four offices. "With a product like Envoy, you can publish a document, and everyone reads it and can add annotations without touching the original."

DAD (Desktop Application Director) and CUI (Common User Interface): DAD gives users quick access to PerfectOffice applications, and users said the CUI helps cut down on training.

Network integration: Users liked the integration of PerfectOffice with NetWare, saying that installation over the network was easy. The GroupWise client provides a single electronic inbox for calendar, E-mail, and schedule items for GroupWise installations.

Some users eagerly anticipate the further network integration that WordPerfect is planning for future PerfectOffice versions. These include integration of GroupWise with NetWare's Directory Services, integration of the InfoCentral calendar with that of GroupWise, and PerfectOffice Select, a CD-ROM of WordPerfect and third-party applications that is due to ship this spring and that will let users select— and pay for—only the applications they need.

—Dave Andrews

GUI Building with Real Intelligence

At first glance, Utah appears to be just another Windows GUI builder, or so I thought when I initially examined the package's documentation and explored the multimedia demonstrations supplied on the CD-ROM. I was skeptical, particularly when the documentation introduced Semantic Information Technology, a term that refers to the model at the heart of Utah. After all, the data-processing landscape is swimming with new "technologies," each poised to "revolutionize" the process of getting a computer to do useful work.

But Utah, which is sold by ViewSoft (Provo, UT, (801) 377-0787), is unique. Yes, it's visual: You click on a button to create a blank window and click on more buttons to populate the window with editable entry fields, scroll bars, radio buttons, or bit maps. Yes, it's object-oriented: The on-screen items just mentioned are not treated as code resources; they are handled as discrete objects that respond to external events (e.g., a mouse-click).

However, Utah goes a step beyond. You don't just build a scroll bar as something to which you will later attach an activity. You can place a scroll bar on-screen only after you've defined a data item that the scroll bar will be linked to. The scroll bar becomes an alternate "view" of the data.

Here's an example: Suppose you had a numeric data-entry field that was a brightness setting for the computer's monitor—zero is completely dark, 100 is full intensity. You might want to provide a horizontal scroll bar in the form of a slider, just under the data-entry field, which the user could adjust with the mouse.

In the Utah way of building GUIs, these two items are alternate views into the same data: one textual, the other graphical. Utah doesn't let you build the data-entry field, build the slider, and *then* decide which data variable to attach them to. It does the opposite: You define a variable and then decide how it is to be represented and manipulated on-screen.

The result is that, in the example I've given, the data-entry field and slider are not only a means of modifying the variable, but they *track* the variable and are thus automatically synchronized with one another. Move the slider, and the number in the data-entry field increments and decrements; modify the content of the data-entry field, and the slider positions itself properly.

Building such objects in Utah requires no coding on the part of the designer. Utah emits the C++ code that you feed to either the Borland or Microsoft compiler. (I tested Utah with Borland's compiler. Utah prebuilds the project for you, ready to drop into Borland's IDE.) Better yet, even before you get to the compile stage, you can test your application in Utah's simulation mode. All the buttons, boxes, sliders, and data fields operate in simulation just as they will in the final application. There's no compile-test-oops-recompile-retest cycle.

The Windows version of Utah costs $495. The NT version costs $695, but that includes the Windows version for free. By the time you read this, it's likely that the Unix version will be available. Utah arrives overflowing with bit maps, buttons, and sliders—everything you need to build a top-notch user interface. And because it's all C++, you can extend the existing set of objects to build custom versions of your own. Utah is definitely worth a look if your GUI building gives you headaches.

THE POWER HUNGRY.

THERE ARE ONLY TWO KINDS OF COMPUTER USERS.

THE TRULY SATISFIED.

CONTENT LICENSING

Legal Land Mines Cloud Multimedia

The world of multimedia development is filled with more hidden legal land mines for developers than any other type of product. The programs bundle many different forms of artistic expression into one product, and each form carries its own standards and practices for getting permission from the creator to make copies. You can't just digitize that neat ending to *Star Wars* and glue it into your game without getting permission from many people.

What it all comes down to when producing a multimedia title are the "two Cs"—clearance and cost—according to Dror Futter, an attorney at the law firm McCarter & English (Newark, NJ). A title such as a multimedia encyclopedia may have 300 video clips, 300 photos or still images, and 300 associated pieces of text, all of which are subject to different licensing and royalty fees.

Clearing the rights to each item and agreeing on a fee to use each piece can be tricky. For example, how do you compare the pricing of a two-page written piece with that of a 2-second sound bite? "Traditional licensing is based on use," says Futter. However, that model can break down with a multimedia application. "You may have to base things on potential use," he says.

Futter gives the following example to illustrate what new types of issues developers of multimedia applications are facing. Suppose you have a virtual-museum multimedia package that displays a set of images every time the program is run (this is analogous to exhibits in the lobby of a real museum). And suppose the package has a Roman art room that requires nine mouse-clicks to get to. "Do we price the art in the Roman gallery the same as the images in the [virtual lobby]?" asks Futter.

Today, there are no clear guidelines to help developers when they face these types of issues. "Right now, everything's done as a custom deal," says Futter. And this has opened the door for specialists who deal in these issues.

For example, Total Clearance (Mill Valley, CA) helps developers wade through the morass of legal requirements, which can be daunting. Jill Alofs, the president of Total Clearance, explains that guilds and unions have set established rates for the reuse of all types of material that will go into film, video, and TV media. But this is not yet the case for multimedia products (e.g., CD-ROM). Multimedia developers will need to consider preexisting content early in their products' life cycles.

—Peter Wayner and Salvatore Salamone

SURVEY: WHAT SHOULD BE COPIED FOR FREE

The U.S. Chamber of Commerce's Telecommunications Infrastructure Task Force surveyed 1600 business users to assess current and planned use of the national information infrastructure. Respondents ranked the types of information that users should be able to copy from the national information infrastructure without fees or restrictions.

1. Anything not copyright-encrypted.
2. Excerpts from articles or publications.
3. Complete articles or publications.
4. Scholarly reports or information.
5. BBS messages.
6. Games, songs, movies, and other entertainment products.
7. Anything they want.
8. Nothing at all.

Current State of Intellectual Property Protection

The U.S. Chamber of Commerce's Telecommunications Infrastructure Task Force surveyed 1600 business users to assess current and planned use of the national information infrastructure. When given a list of possible amendments to existing copyright laws, participants of the U.S. Chamber of Commerce's survey generally favored amendments that clarify when information can be downloaded or transmitted to third parties.

	Define the types of information that can be downloaded without fees or restrictions.	Define the types of information that can be downloaded with added fees and/or restrictions.	Allow copyright holders to restrict or charge extra when the transaction is equivalent to a sale (e.g., downloading a document or publication).	Provide the business/consumer and the service provider the conditions under which downloading and/or transmitting materials to third parties will be allowed.	Prevent users from transmitting materials downloaded from the national information infrastructure to third parties under all circumstances.
Favor	79%	70%	71%	78%	23%
No opinion	10%	10%	12%	12%	13%
Oppose	11%	20%	7%	10%	64%

BYTE

EXTRA

International

Edition

◀ **Illegal Interference**

Starting next year, European offices will be safer. The European Community's new directive will protect computers and other electronic equipment from nasty electromagnetic disturbances.
page 32IS 3

Scanner/Digital Camera
page 32IS 14

24-Bit Full-Motion Multimedia System
page 32IS 13

Networking for Satellite Offices
page 32IS 18

THE SMART ALTERNATIVE..

Snap into the future with the LEO DESIGNote which provides full flexibility and upgradeability in a modular package. The interchangeable LCD allows effortless upgrades from a mono STN display to either a color dual-scan STN or color TFT display. The removable HDD lets you swap drives with up to 500 MB of storage, and you can also replace the 3.5" FDD with additional modules.Intergrated with VESA local bus VGA Graphics, a built-in trackball, a built-in microphone and speaker for on-board audio, ports for PCMCIA type II or III, a choice from a full line of CPUs, and a host of other innovative features, the LEO DESIGNote is the most powerful and versatile notebook solution around.

intel inside ®

... *LEO DESIGNote*

Illegal Interference

ADELE HARS

RICHARD STEADHAM © 1995

Your report's due at noon. You need to just print out that one last page and ... the network goes down. Everyone scrambles around to pinpoint the cause, but they find nothing. Who would think that it's because someone in the next building turned on the air conditioning? When they hit that AC button, it created a power surge that radiated an electromagnetic field across the network cabling.

"If something goes wrong with your computer because of something outside of your visual range, you aren't aware of it. All you know is that something's gone wrong with your computer," says Rob St. John-James, project engineer at IBM's EMC Laboratory in the U.K.

But by January 1996, every system or piece of computer equipment (and all other electric or electronic products destined for end users) you buy in Europe has to be stamped with a "CE." That should mean—with certain caveats— that under normal circumstances you need not worry about EMC (electromagnetic compatibility) concerns.

The CE marking indicates the manufacturer has determined that the product will neither cause nor be susceptible to electromagnetic emissions. Most manufacturers will use the services of an officially recognized test house to achieve this. (The CE marking will also signify compliance with other standards, but for now, the big concern is EMC.)

Most countries have some kind of EMC legislation, with Germany's being regarded as the strictest. But everyone (including the U.S. and Japan) has dealt only with the amount of electromagnetic emissions a piece of equipment can give off. What Europe is doing is unique: It is requiring manufacturers to also prove that their products are sufficiently *immune* to common forms of electromagnetic radiation, such as those that might come from power transients, TV and radio broadcasts, electrostatic shock, lightning, and other causes.

Europe's upcoming EMC regulations should help shield equipment and buyers from electromagnetic disturbances

Designing and testing products for electromagnetic immunity has sent many manufacturers back to the drawing board, sometimes for major redesigns. One engineer estimates that without at least some redesign, 90 percent of products tested would fail the new EMC tests. Everything has to be considered: the chips themselves, how they're placed on the board, the number of planes on the board, how the boards are placed inside the system, the location of the power supply, and cabling. Furthermore, new products are coming out from companies like Siemens Nixdorf and AT&T to address cabling for LANs, a big source of EMC concern.

Awareness campaigns are rolling out across Europe. The U.K.'s DTI (Department of Trade and Industry) has conducted an EMC Update campaign for several years to let manufacturers know what they need to do to comply. The French Ministry of Industry, on the other hand, is just getting started with a similar effort.

In Germany, says Allen Emery of the Siemens Nixdorf test lab

What Causes Noise in the Office?

High-current power cable

Network cable

Main socket

Video cable

A typical office can be a hotbed of electromagnetic disturbance. According to Interference Technology International, the scene depicted here represents the emissions problems found in a normal IT environment. The computer interferes with the radio because of radiation from printed circuit boards within the inadequately shielded computer case, as well as from the various cables. Emissions from the high-current power cable can interfere with the network wiring. The mobile phone's transmissions can interfere with the PC. And incoming transients on the main power line affect internal power supplies due to inadequate filtering.

in Paderborn, Germany, the big companies are prepared, because they typically have people working on the standardization committees anyway. "Interest groups are being formed by smaller companies to gain information regarding the directive and testing facilities," he adds.

Retailers are also reacting. Many major U.K. outlets are saying that after this June they will no longer accept delivery on products that do not carry the CE marking, ac-

cording to Tony Maddocks of ERA Technologies, a DTI-approved consultancy specializing in EMC issues.

In the meantime, many users are unaware of the directive and its potential impact on what they're buying. Michel Lambert, who coordinates network security for the French insurance giant Groupama, says that while he has heard of the directive, it is not yet a big issue. On the other hand, Gianclaudio Puglisi, IT (information technology) direc-

tor for PSA Peugeot Citroën, says that while his company does not yet include EMC specifications in its purchasing requirements, it's an issue that is "very, very important." The industry must respect such environmental concern, he adds.

Dave Southerland, who is responsible for implementation of the EMC directive at the U.K.'s DTI, says his organization has been contacted by banks, insurance companies, research and medical institutions, and even other government agencies asking what they should be purchasing. Siemens Nixdorf and IBM confirm that they have had numerous customer inquiries, too. "There is a growing awareness because more and more people are seeing the CE mark," says Southerland.

It Will Make Life Easier

Theo Dirksen, customer assurance manager for Hewlett-Packard's PC division in Grenoble, France, points out that for consumers, the benefit of the CE marking is that they'll be able to buy products from other European countries without worrying that they don't conform with standards in their own countries. Transnational corporations that have centralized purchasing will have to worry about only one mark (CE) rather than 15, he adds.

This will be a welcome change for IT

TIPS FOR AVOIDING EMISSION PROBLEMS

The EMC (electromagnetic compatibility) directive will cover most people in most situations. This should not be an excuse to throw caution to the wind, however. Rob St. John-James of the IBM EMC Laboratory in the U.K. offers users the following advice for avoiding EMC problems:

- Make sure your equipment and cabling meets the EMC standards.
- Use properly screened cables or I/O filters; unprotected or badly installed network cabling can effectively act as

a big EMC antenna.
- Don't allow radios, mobile phones, paging systems, and walkie-talkies in the computer room.
- Review your site for potential EMC culprits.
- Make sure you've got a clean power supply.
- Read your user's manual, and follow the manufacturer's installation and maintenance instructions.
- Before you perform any major IT (information technology) installation, get a professional EMC site survey.

managers who work for pan-European businesses, such as Adrian Luther, hardware planning manager for banking giant Barclays. He specifies computer purchases for use at sites across Europe. "It will make life easier," he says. He can also move equipment around from country to country without worrying about violating national standards.

Dave Haddleton, who oversees computer purchasing, says that Barclays is well aware of the EMC problem, because it once had to shield certain personal computers that were affected by nearby high-frequency cables.

David Mylrea, EMC engineer with Digital Equipment in Scotland, points out that the Alpha processor has a lot more transistors than most chips and draws more current. Unfettered, it could have been an EMC nightmare (generally, the faster the chip, the greater the emissions), Mylrea says. But Digital built in EMC considerations from the beginning of the chip design process. As a result, says Mylrea, "It's a very quiet chip."

Look for this mark when buying electronic equipment. It certifies compliance with new EMC regulations.

But what will this do to price? Whether or not you see increases will vary from product to product, depending on how much redesign is required. The big players typically can absorb the costs without passing them along to consumers. For the small players, it's more difficult. On the other hand, because companies can test for one group of European standards instead of on a country-by-country basis, there should be some savings.

Who Is Responsible?

Generally the manufacturer is responsible for ensuring that equipment meets the new standards, but not always. The buyer also has certain responsibilities. Tony Maddocks of ERA Technologies stresses that a user must operate any piece of equipment under the conditions the manufacturer specifies in the documentation. If a user fails to follow instructions (e.g., using an add-in card that the manufacturer doesn't list as compatible) the manufacturer will not be held responsible for any EMC problems. Furthermore, just because your screen flickers doesn't mean you have a court case against a manufacturer. The important thing is that data integrity is maintained.

There has been debate regarding whether PC add-in boards aimed at the end user need to be CE-marked. While some have argued that the boards alone have no "intrinsic" purpose (one of the criteria cited in the directive), the consensus appears to be that if the end user buys it—at least in a mass-market context—it should have to carry the CE mark.

The final responsibility for compliance with the EMC directive lies with whoever puts the equipment on the market, explains IBM's St. John-James. This is not just a problem for manufacturers. Mail-order companies, systems integrators, VARs, OEMs, stores, and even systems administrators assembling a system for their own company may ultimately be responsible. This question of who is ultimately responsible has not yet been put to any legal test. But according to the way the directive is worded, whoever puts a computer system together is responsible. For example, if a reseller configures the system, that reseller is responsible.

It's worth noting that the directive differentiates between systems and installations: Systems must be branded with the CE mark; in installations, the component parts must be marked. However, will dealers who decide to run a special—taking a CPU from one manufacturer, a monitor from another, a printer from a third, and a modem from a fourth—be responsible for EMC conformity? Possibly, if the package deal is sold as a

# Ergonomic Furniture for every Computerized Work Environment

ORGANIZE COMPUTER EQUIPMENT

Ergotron LAN Furniture organizes monitors, servers and other equipment upward into previously unused "free" air space. For existing equipment , this immediately reduces floor space requirements. As your LAN grows, Ergotron can help you potentially triple or even quadruple available floor space without costly building expansion.

IMPROVE COMPUTER ACCESS

With the flexibility of Ergotron's open architecture, you can custom-configure standard components to organize equipment, improve computer access and reduce system downtime. The unique Cable Management System not only hides cables, but organizes them for easy access and identification.

ERGONOMIC ADJUSTABILITY

And with Ergotron's patented monitor suspension, you can lift, glide, tilt, swing, and rotate every monitor until they're "just right" for easy access and comfortable viewing.

Tested at over 900kg, Ergotron systems are ruggedly built to protect your computer investment and last a lifetime.

For more information and a free brochure, contact our European sales office:

Ergotron Europe
Karspeldreef 19.1102 BB
AMSTERDAM
The Netherlands

Tel: +31 20 696.60.65
Fax: +31 20 609.04.59

ERGOTRON EUROPE
Ergonomic Computer WorkCenters

"system." The same thing goes for systems integrators. But again, none of this has yet been put through a court challenge.

It's still very much a case of "buyer beware." Just because a product is CE-marked does not mean it is invincible. It is your responsibility to read the manufacturer's installation and operating instructions—and follow them.

When you buy a board and put it into your computer, you assume a certain degree of responsibility. There's no way that a PC manufacturer, for example, can test for EMC with every possible plug-in board on the market. If you plug in a board that the PC manufacturer has not cited in the documentation, you may assume liability for any EMC problems that might arise once the board is installed. If your neighbor complains that your PC, to which you've added a board, is causing EMC problems, you could be obliged to replace the board.

That's why you should only buy boards that are CE-marked, says Robert Muir, senior design development engineer with Inmac, a mail-order house in the U.K. Inmac still has a small share of the networking-products market, he says, and it's an area where price competition is tight. This means companies like his can't pass along the EMC product design and testing costs (which may account for up to 30 percent of the production cost) to their customers. Muir says he is concerned that unscrupulous manufacturers may stick a CE mark on their products without testing them properly.

Policing is a big issue. Ultimately, it is up to consumers to speak up if they have an EMC problem and file a complaint with the national authorities. The problem is that you often have to be Sherlock Holmes to track EMC problems down.

Unless it occurs to someone that, for example, the network goes down at the same time each day when a nearby radio transmitter starts broadcasting, many EMC problems go undetected. St. John-James estimates that perhaps as little as 2 percent of EMC problems are ever reported.

The problem will become even more acute with the mushrooming of mobile phone use and other remote communications devices, which are well-known EMC culprits, and with the ever-increasing speed of computer chips. Other EMC troublemakers are fluorescent lights, TV and radio broadcast towers, power lines, electric motors, lightning, and power that is not properly grounded.

Sins of Emission

- A surgeon in France reported that during an operation his medical instruments started displaying erroneous information due to an electromagnetic disturbance generated nearby.

- A major West End show was delayed many weeks because the computer-controlled scenery would unexpectedly move each time someone on the street made a call with a mobile phone.

- One company lost its telephone system and many modem cards when lightning struck nearby. The EMC (electromagnetic compatibility) directive covers such surges. If the cards had been CE-tested, says Peter Brown of Interference Technology International, they might have been protected.

Despite Questions, the Right Move
The EMC directive relies on European and international standards to spell out the details, and most of these standards are not

Checking a computer for CE-compliance in a chamber at IBM's EMC test lab in the U.K. The lab is a DTI-appointed Competent Body.

yet in place. While the general standards will be ready for 1996, many of the product-specific ones will not be.

The large manufacturers, who are often members of the standards bodies, generally know what they need to do to comply. However, small manufacturers are often hampered by a lack of reliable information. Regardless, it's up to each company to find out what they are required to do to comply with the new regulations. Companies that aren't certain what they are required to do should contact their local chamber of commerce or national standards organization.

All these electromagnetic problems won't miraculously disappear, but the EMC directive addresses a problem that, until now, has been almost literally invisible. As Philippe Brawerman, president of Cisco Systems Europe points out, it's a step in the right direction for manufacturers selling across Europe. Because they'll have just one certification to obtain for EMC instead of 15, maybe we'll see technology and products getting out a little faster. ■

Adele Hars is a BYTE correspondent and a journalist with the GEID Press Agency in Paris. You can reach her via CompuServe by sending E-mail to 100325,3703 or on BIX c/o "editors."

For More Information

ERA Technologies
Leatherhead, Surrey, U.K.
+44 1372 367068
fax: +44 1372 374074
Circle 976 on Inquiry Card.

Interference Technology International
Shrivenham, Swindon, U.K.
+44 1793 783137
fax: +44 1793 782310
Circle 977 on Inquiry Card.

OBJECT.STATION 41

Canon Computer Systems' object.station 41 is a system based on an Intel 100-MHz 486DX4 CPU that's built to optimize NextStep. At the heart of the object.station 41 is a Canon video subsystem that speeds video throughput at NextStep's native resolution while providing compatibility with most video standards. Available from P&L Systems (Old Amersham, U.K.), the system matches NextStep's designated workstation-class video resolution of 1120 by 832 pixels and supports the standard video resolutions required for Windows, DOS, and other platforms.

System features include 2 MB of video memory, which allows video resolutions of up to 1280 by 1024 pixels, and 16 MB of RAM, which you can expand to 96 MB. Other features include CD-quality audio with Sound Blaster capability, a 32-bit LAN interface, SCSI-2 controllers, three ISA slots, one VL/ISA slot, a speaker and microphone, a 17-inch color monitor, a 500-MB hard drive, and keyboard-resident volume and brightness controls.
Price: £3333
Phone: +44 1494 432422 ;
fax: +44 1494 432478.
Circle 1023 on Inquiry Card.

TWO-SLOT VMEBUS PCMCIA CARRIER ▼

Arcom Control Systems (Cambridge, U.K.) has added a two-

slot PCMCIA carrier for Type I, II, and III cards to its VMEbus range. The VPCMCIA lets you install, run, and upgrade standard DOS and Windows applications in remote locations. Its dual slot accepts most PCMCIA cards, so you can combine storage devices with communications links via fax modems or LAN interfaces.

Designed around the Cirrus Logic CL-PD6720 controller, the VPCMCIA includes a 128-KB

CORDLESS PC CONTROLLER

A cordless pointing device, TrackMan Live is optimized for use during computer-based presentations. Featuring cordless radio control with a range of up to 30 feet, the device lets you focus on your audience and circulate freely while speaking.

The device also functions as a thumb-operated three-button mouse, so you can edit your work up to the last minute and even during the presentation itself. Enhanced MouseWare software, which ships with the unit, permits customized commands for the second and third buttons and additional personalization of cursor size and sensitivity.

TrackMan Live features four channel options to prevent interference with similar devices in use nearby; sleep mode; a low-range battery option, which is an alternate 6-foot range to conserve energy and prolong battery life; and a battery status signal, which indicates if you have adequate power.
Price: £169
Contact: Logi Ltd., Windsor, U.K., +44 1344 894300;

EPROM socket for local storage of configuration data, and a DC power-distribution architecture supporting 3.3- and 5-Vcards. Extra logic circuitry on the card lets any Arcom VMEbus system support up to two carrier cards or four PCMCIA slots. Each PCMCIA socket has five programmable memory windows of up to 16 MB in size.
Price: £595
Phone: +44 1223 411200;
fax: +44 1223 410457.
Circle 1029 on Inquiry Card.

LOW-COST FRAME-RELAY ACCESS

The Vanguard FRAD (Frame-Relay Access Device) for LANs and SNA networks provides SLIP-to-TCP/IP conversion with RFC 1490 encapsulation, SDLC-to-LLC2 conversion using RFC

1490, SDLC-to-QLLC conversion for interfacing to X.25 networks, and support for SNA, Burroughs Poll/Select, BSC3270, BSC3780, and 25 other protocols. In addition, the device includes virtual PU mapping, SNMP support, and an optional second port for backup or additional protocol support. With Motorola's (Crawley, U.K.) Vanguard FRAD, you can remotely download software as requirements evolve.
Price: £560
Phone: +44 1293 404343;
fax: +44 1293 404362.
Circle 1024 on Inquiry Card.

INTELLIGENT LINE BOOSTERS

The DataBoost parallel and serial line drivers offer error detection and correction. Each of the DataBoost modules incorporates a high-speed microprocessor, which assembles the data into packets before transmission. The packets contain a CRC-16 checksum, which lets the receiving module detect errors in the received data. When a module detects any errors, it retransmits the data. The full DataBoost set includes both transmitter and receiver modules.

The modules from R-Tek Research (Luton, U.K.) include DataBoost iP, a high-speed parallel line driver; DataBoost iPS, an auto-converting parallel-serial driver; and DataBoost iS, a full-duplex RS-232 line driver with flow control. The parallel and serial modules are fully compatible, so you can use parallel-to-serial or serial-to-parallel pairs. When you use them together, the serial modules allow fully bidirectional data flow over a single twisted-pair cable with full flow control at both ends of the link.
Price: £149
Phone: +44 1582 490092;
fax: +44 1582 490046.
Circle 1026 on Inquiry Card.

Hardware

SPEECH-RECOGNITION ▶ SYSTEM

ALPSpeak for Windows, the speech-recognition system from ALPS Electric Ireland (Millstreet Town, Ireland), is a speaker-independent system that plugs into your PC's printer port. The unit works with popular Windows applications, such as Word, Excel, PowerPoint, and Access.
Price: about £299
Phone: +353 29 70677;
fax: +353 29 70603.
Circle 1031 on Inquiry Card.

ELF MONITOR

With Perspective Scientific's (London, U.K.) ELF Monitor, you can detect and measure electromagnetic radiation levels in your workplace and home. The palm-size device can measure radiation over the range of 20 to 1200 Hz. The ELF Monitor senses the magnetic field by a pickup coil, which produces a voltage that it amplifies and processes to drive the LED

screen. The sensor has a loop area of 30 square centimeters, and the screen has 10 LED lights, each denoting a level of exposure from 1 to 24 milligauss. The LEDs are grouped in ranges of exposure: probably safe levels (200–300 nanotesla), in excess of safe levels (600–800 nanotesla), and warranting health concerns (above 8 milligauss).
Price: £110
Phone: +44 171 486 6837;
fax: +44 171 487 3023.
Circle 1025 on Inquiry Card.

24-BIT FULL-MOTION MULTIMEDIA SYSTEM

Highmead Technologies' (Welham Green, U.K.) Fountain FM1 Movie Master System is a multimedia system based on a 66-MHz 486DX2 CPU that has a 24-bit MPEG card, which provides full-screen, full-motion video with 24-bit resolution and high-quality digitized sound. The system comes with 4 MB of RAM, a dual-speed CD-ROM drive, a 16-bit sound card, and three video CDs. The MPEG card is fully compatible with the MPEG 1 video format at all frame rates and features full compatibility with the MPEG audio format in all modes. The Fountain FM1 Movie Master System boasts playback of video-CD movie and music discs, plus the ability to display still images with full scalability at the card's output resolution.
Price: £1249.95
Phone: +44 1707 272627;
fax: +44 1707 267400.
Circle 1030 on Inquiry Card.

HIGH-SPEED COLOR SCANNER

The 600-dpi JX-330 can scan up to 16.7 million colors with 256 smooth shades of gray. Using Sharp Electronics' proprietary algorithm, it is capable of producing scans of up to 2400 dpi. Both PC- and Mac-compatible, the JX-330 offers you a choice of two film-scan units, each fitted with a CCD reading unit, and provides the ability to produce high-quality scans of slides and transparent art. These options allow you to handle film sizes of 35mm up to 211.3mm by 279.4mm. The first film-scan unit (JX-3F6) provides resolutions of up to 1200 dpi, while the second one (JX-3F12) goes up to 2400 dpi. The JX-330 has an optional 50-page automatic document feeder, and you can install it and the transparency unit at the same time.

Price: £1395
Contact: Sharp Electronics Ltd., Newton Heath, U.K., +44 161 205 4255;
fax: +44 161 205 2638.
Circle 1021 on Inquiry Card.

Industrial Strength Applications!

```
WIN-PROLOG - [Console]
File   Edit   Search   Run   Options   Window

Build industrial strength applications for the industry
standard platform, Windows 3.1! Robust, intelligent
applications which yield the full 32-bit power of your PC.

And with LPA-PROLOG, build them the easy way:

    High level handling of dialogs, menus and graphics
    Powerful DLL, DDE and ODBC interfaces
    Compact, fast and royalty-free runtime system

Not to mention the integrated source level debugger,
multi file program editor, incremental and optimising
compilers, and full Prolog predicate library.
```

LPA WIN-PROLOG
The industrial strength compiler for business applications!

LPA Logic Programming Associates Ltd
Phone (US Toll Free): 1-800-949-7567
Phone: +44 181 871 2016 - Fax: +44 181 874 0449
Email: lpa@cix.compulink.co.uk - Compuserve: 100135,134

Highspeed Data Communication

T.U.S.S.I
460'800 bps

- Ultraspeed Serial Card up to 460'800 bps
- DSP-technology/ISA 16 bit card
- 100% register comp. with 16550 UART
- On-board 32'000 byte buffer
- Independent from Operating System
- No special software drivers needed

SAFELine

- Data Encryption TDE
- V.34, update to ISDN
- ECM-Fax Error Correction
- Leasedline-Interface
- Modular
- Standalone/Rack/PCMCIA Version

SWISSMOD

- V.34/V.32bis
- Data/Fax/Voice
- Build-in Fax/Data/Voice-Switch
- Remote-Switch
- Conversation Recording
- Standalone/PCMCIA Version

Distributors and Dealers wanted!
OEM versions available
Info-Service D 06172-21096

TELELINK
DIE SCHWEIZER **MODEM**MACHER

Gewerbestr.11, 6330 Cham
Switzerland
Phone: (41) 42-43 10 80
Fax: (41) 42-43 10 81

Hardware

SCANNER/DIGITAL CAMERA

The Lumina scanner/digital camera captures production-quality reflective, transmissive, and 3-D images of still objects. The camera accepts Nikon bayonet-mount SLR lenses, and it can use traditional creative photographic accessories such as filters and lighting.

A single-pass, rapid-scanning camera, the Lumina houses a 12-bit trilinear CCD array that moves rapidly through 3400 steps. This produces an image with a density rating of 3.0D at 600 dpi or 300 dpi with A4 output. Maximum resolution is 2400 dpi via software interpolation. You can produce a full-screen preview at 72 dpi on a 19-inch monitor in just 20 seconds.

The Lumina comes with a SCSI-2-compatible interface plus two Photoshop plug-ins, which allow custom toning of files. You can save files in a 24-bit format for Photoshop manipulation or as a High Dynamic Range format for archiving, eliminating the need to rescan.

Price: £5995
Contact: IMC, Windsor, U.K., +44 1753 830999; fax: +44 1753 831094.
Circle 1022 on Inquiry Card.

INTERNET IN A BOX

To simplify access to the Internet and other on-line services, Computers Unlimited (London, U.K.) offers Internet In A Box, a complete communications solution consisting of a modem, software, cables, and reference books. The package includes the MultiTech MultiModem ZDX 19.2-Kbps data/fax modem, the Internet Starter Kit for Macintosh, the CompuServe Mac Membership Kit, Snatcher file transfer software, and MicroPhone LT, an entry-level communications package from Software Ventures.
Price: £299
Phone: +44 181 200 8282; fax: +44 181 200 3788.
Circle 1036 on Inquiry Card.

MESSAGE SERVER

Isoplex V4 is software that combines rapid message-switching ability with a commercial-grade scalable message server that can support thousands of users. The message server facilitates the seamless exchange of information, including text, forms, spreadsheets, and multimedia objects.

Isocor's (London, U.K.) multiprotocol message server offers accounting support, advanced security features, directory services based on X.500, local and remote management, gateways to popular E-mail systems, EDI using X.435, faxing, and multimedia file support.

Isoplex V4 features alternate-path routing, rules-based distribution, and auto-forwarding. A Message Store option supports client/server access from dial-up and LAN-based users on a multiport machine. Users on Windows PCs or Macs can dial directly into Isoplex V4's Message Store using the X.400 P7 protocol, which operates transparently under MAPI or System 6 or 7.
Price: £5600 to £75,000
Phone: +44 181 754 8115; fax: +44 181 754 8120.
Circle 1028 on Inquiry Card

MT printhead
refurbishment

portable printers

4-40 ppm lasers

Mannesmann
Tally
delivers printer
SOLUTIONS
to match your
NEEDS

160-900cps
matrix

barcodes, labels,
ticket printing

Heavy-duty
line printers

MT flexible
maintenance

Colour ink-jets

MANNESMANN

Tally **It pays to think Mannesmann Tally**

Call +44 1734 788711 for details of your local distributor or fax +44 1734 771312

T

THE STRONGEST PRINTER RANGE IN THE WORLD

Circle 533 on Inquiry Card.

PRINTER MANAGEMENT ▲

The latest version of Interlink Communications' (Clacton-on-Sea, U.K.) Netprint range of LAN print servers, Netprint PDS is designed for use in structured-cabling systems. The device requires only a single UTP or thin-wire Ethernet LAN connection for support of up to 16 printers. Connection to the structured cabling is via a series of 15 RJ-45 connectors similar in concept to a LAN hub. Printer connection is via Netpod modules that plug directly into the printer's parallel or serial port. Netpod connection to the structured cabling is via an integral 5-meter fine-profile ca-ble terminated with an RJ-45 plug. You can daisy chain Net-pods together, connecting a num-ber of printers from a single desk or wall point. Netprint PDS also supports the direct connection of a parallel printer via a high-speed bidirectional printer port.
Price: Netprint PDS, £395; Net-pod modules, £95 each
Phone: +44 1255 426147;
fax: +44 1255 221369.
Circle 1038 on Inquiry Card.

BRIDGING THE TRANSPUTER GENERATION GAP

The QTC100, a hardware mod-ule, provides a communications bridge between existing T400- and T800-based systems and the new-generation T9000 transputer. A size-2 TRAM, the QTC100 converts the four O/S links pro-vided through the TRAM inter-face into differential T9000 D/S links, which it makes available through standard connectors. Other features of the Quintek (Westbury-on-Trym, U.K.) mod-ule include selectable link speed, reset and control capability, com-patibility with existing develop-ment tools, and existing system host capability.
Price: £750
Phone: +44 1272 628196;
fax: +44 1272 628717.
Circle 1027 on Inquiry Card.

LAN CARD

A new-generation low-cost net-work card is compatible with Ethernet networks. Nighthawk Electronics (Saffron Walden, U.K.) offers two versions of the card, BNC only or ComboPlus, in-corpo-rating three on-board trans-ceivers for UTP and thin-wire coaxial cabling. Both versions are NE2000-compatible, provide a 16-KB RAM buffer, and have a 10-Mbps transmission rate. A socket is also included for optional re-mote-boot EPROM. The network card does not have any DIP switches or jumpers, making in-stallation simple. The menu-driv-en software covers interrupts, I/O, and boot-ROM addresses.
Price: BNC only, £41; Combo-Plus, £48
Phone: +44 1799 540881;
fax: +44 1799 541713.
Circle 1037 on Inquiry Card.

THE MOUSE IS DEAD.

LONG LIVE THE MOUSE!

But the general input and control device of the future must have more to offer than three buttons: More buttons, freely assignable. Senses enough to read and size graphics. To digitize anywhere in WINDOWS, even where digitizer tablets are not wanted, in CorelDRAW! and AutoSketch for instance. Senses to support menu templates everywhere and to navigate you safely through large spreadsheets. Everything the mouse had done plus more than all graphics tablets can do. Without any fuss. This is the future. It's here now.

"It's a technical marvel." **Windows Magazine 6/94** "Extremely versatile and easy to use." **PC Home 5/94** "They said it couldn't be done. Digitising within AutoCAD LT was supposed to be out of bounds. ARISTO, the German manufacturer of digitisers, begs to differ. And, what is more, have a product to prove the point." **CADDesk 5/94** "Multitalent für Windows-Freaks" **AutoCAD 1/95**

THE 𝐴𝑅𝐼𝑆𝑇𝑂 MOUSE.

New Technology – New Opportunities

Circle 534 on Inquiry Card (RESELLERS: 535).

ConEQ/DB v2.2.4

CONEQ TO QUALITY

ConEQ/DB is a powerful, simple to use, graphical and statistical program for SPC with database management, running in DOS or any PC network. Major users are certified suppliers of GM, FORD, RENAULT, PEUGEOT and others. **They are ConEQted to quality**. Don't settle for less. Free demo available at our anonymous FTP server at *asterix. intersis. pt*

INTERSIS S.A.

INTERSIS S.A. Est. Paço d'Arcos, 48 • 2780 OEIRAS •PORTUGAL
Tel. 351.1.4412247 Fax 351.1.4412287 E-mail: info@intersis.pt

---

What's New

SPACE-SAVING TFT COLOR LCD MONITOR

Sharp Electronics (Newton Heath, U.K.) offers the LC-10C1M/U, a 26-cm diagonal TFT color LCD monitor that combines high-definition images and vivid color reproduction. Using a low-reflection, flicker-free LCD panel, the LC-10C1M/U displays up to 4096 colors, incorporates 307,200 pixels, and is compatible with VGA, EGA, and Apple Mac II series modes. The LC-10C1M/U display reads the input signals and automatically switches the VGA modes and Mac II modems by an analog/digital changeover switch that's on the back of the unit.
Price: call for price
Phone: +44 161 205 4255; fax: +44 161 205 2638.
Circle 1035 on Inquiry Card.

NETWORKING FOR SATELLITE OFFICES

If your company has sites where provisions for remote access couldn't previously be justified, SpiderIntegrator Atto may be of interest to you. Spider Systems (West Drayton, U.K.) designed the entry-level internetworking product to provide fixed and on-demand remote access for small satellite offices.

The device provides direct connection to leased lines or ISDN and supports synchronous dial-up. You can link sites using SpiderIntegrator Atto to SpiderIntegrator Mezza and Pico units at a central site and manage them from that central site. Atto can also be used to connect small offices together without the need for another router.

SpiderIntegrator Atto is available in two configurations: a remote bridge that transparently bridges all protocols, and a remote router, capable of routing IP and Novell IPX protocols. In its remote-bridge configuration, the product supports one LAN and two serial ports or one serial port and one ISDN basic-rate interface. As a remote router, Atto supports one LAN, one serial port, and one ISDN BRI.
Price: remote bridge, £995; remote router, call for price
Phone: +44 131 554 9424; fax: +44 131 554 0649.
Circle 1032 on Inquiry Card.

MULTIMEDIA INFORMATION POINT

A product called the PIM can supply information, certifications, bookings, and ticket services. The device's external structure is made by ABS in a choice of colors. Inside the unit is a chassis of laser-cut galvanized sheet iron with a forced ventilation system, thermostat, and magnetothermic main switch.

The intelligent section of the PIM, from Ecosystems (Florence, Italy), is based on a 66-MHz 486DX2 CPU with 128 KB of cache memory, 4 MB of RAM, a 210-MB hard drive, a 1.44-MB floppy drive, a 16-million color SVGA board, a 14-inch color monitor, a touchscreen infrared system, and a sound board linked to a stereo-output hi-fi system. Options include powerful processors, CD-ROM readers, optical disks, printers, 32-column check printouts, money-collection devices, money readers, and badge readers.
Price: about 9 million lira
Phone: +39 55 353966; fax: +39 55 354306.
Circle 1039 on Inquiry Card.

ISDN FOR THE MAC

The new iSDN Manager from 4-Sight (Bournemouth, U.K.) lets you rapidly transfer files between locations and run third-party communications packages at high speed. Its Macintosh Connection tool lets communications packages automatically redirect their traffic over the ISDN link and lets iSDN Manager offer a WAN system for AppleTalk networks. With a package such as Apple Remote Access, you can link distant networks via an ISDN connection. Once connected, mail systems, file servers, printers, and databases become transparently available to the remote user.

Version 2.0 includes a Quark Xtension, which automatically sends documents with associated graphics files from within Quark, and an Adobe Acrobat Plug-in, which lets Acrobat and iSDN Manager users automate file transfer for instant client-proofing from within Acrobat. It also features Desktop ISDN Out baskets for sending files by simply dragging the files or folders to an ISDN icon on the desktop.
Price: with card, £1895
Phone: +44 1202 764401;
fax: +44 1202 761666
Circle 1050 on Inquiry Card.

ELECTRONIC POINT-OF-SALE SOFTWARE ▼

According to Automated Programming Techniques (St. Albans, U.K.), ShopKeeper keeps costs down by reducing the need for support. The package provides facilities for scanning, descriptive billing, discounting, velocity codes, account sales, deposits, cashier control, and stock. In ad-

ALL-IN-ONE SOLUTION

RVS-COM for Windows 2.0 combines functions such as file transfer, terminal emulations, videotext, and a mailbox, as well as sending and receiving faxes, E-mail, and voice mail. The package supports various data transfer methods, modems, X.25 cards, and ISDN-Common-API cards.

The DEC-VT and European standard videotext terminal emulations display all colors, graphics characters, attributes, and display formats. You can choose any size terminal window. The program's MultiReceiver can receive faxes, E-mail, data files, and voice mail. It automatically detects the type of incoming call. Voice-mail functions let you receive, record, play back, and convert voice messages.

With the mailbox, you can run a multiline BBS with different user levels and individual user accounts. The MDI application lets you use the address book, text editor, terminal, and Combas program simultaneously. DDE capabilities let you exchange data with other Windows applications.
Price: DM 490 to DM 1590

Contact: RVS Datentechnik GmbH, Munich, Germany, +49 89 3571570; fax: +49 89 35715799.
Circle 1042 on Inquiry Card.

dition, ShopKeeper provides interfaces for EPOS devices, such as credit-card readers, bar code scanners, weigh scales, cash drawers, receipt and slip printers, and customer displays. At any time, you can add optional features, such as promotions and EFTPOS. ShopKeeper is available in American English, Dutch, Finnish, French, and German. The package can handle multiple currencies and languages. The source code is available if you want to tailor your own system.
Price: from £150
Phone: +44 1727 835155;
fax: +44 1727 839881.
Circle 1044 on Inquiry Card.

CLASSIC COLOR CLIPS

If you want to add impact to your documents, memos, newsletters, presentations, catalogs, reports, and brochures, SoftKey International's (London, U.K.) PC Paintbrush ClipArt Library may interest you. The library of 12,000 images on CD-ROM covers 161 categories, in-

cluding nature, sports, and technology. The images are scalable, vector line-art images, so you can increase or decrease the size and aspect ratio without reducing image quality.
Price: £46.95
Phone: +44 181 789 2000;
fax: +44 181 789 5626.
Circle 1046 on Inquiry Card.

STORE AND RETRIEVE ▶ PRESS CLIPPINGS

With ClipNote, you can scan pages containing interesting items into a Lotus Notes system. You can then crop the scanned pages and remove uninteresting text, leaving a clear representation of the relevant item. You can also highlight the relevant area. ClipNote offers annotation, letting authorized users add subjective comments that may be read by other users.

The system from Headway Image Enabling Group (Aldershot, U.K.) uses a flatbed image scanner to input images to ClipNote from single pages or magazines and company reports. After scan-

ning, you get preprocessing options, such as automatic separation between text and graphics information on a page, de-skewing, and de-speckling. ClipNote stores scanned and indexed images on optical disks, providing full network access.
Price: call for price
Phone: +44 1252 333575;
fax: +44 1252 314445.
Circle 1058 on Inquiry Card.

Making headway

What's New

PROJECT MANAGEMENT ▲

Previously available for the Mac, Micro Planner X-Pert is now available in a Windows version. Companies with a multiplatform environment can share project information transparently between these two operating systems. Micro Planning International's (Westbury-on-Trym, U.K.) program lets you build your network on-screen, adding and rearranging tasks as necessary.

Micro Planner X-Pert breaks each phase of a project into its component tasks, so that all concerned can see what

has to be done when, in what order, and by whom. Once a project is under way, you can update the plan and indicate areas requiring special attention. Flexible reporting lets you produce personalized schedules for key individuals and subcontractors.
Price: £1750
Phone: +44 1272 509417;
fax: +44 1272 508377.
Circle 1057 on Inquiry Card.

ARCSERVE IN SPANISH AND ITALIAN

In addition to its native French and German versions, Cheyenne Software (Red Hill, U.K.) now offers native Italian and Spanish versions of ARCserve for NetWare, Windows Edition. The NetWare backup program features AutoPilot, for automatic tape rotation and disaster recovery, and Media Management, for on-line database backup.

SIMULATION GENERATOR FOR THE PC

Nonprogrammers can quickly create interactive simulations and prototypes on-screen
with Rapid. You can use the software tool to design human-machine interface systems such as cellular phones and medical equipment.

The Rapid development system consists of two main modules. You use the Object Editor to define the application's building blocks by selecting and customizing controls, switches, and displays from an object library. Using a visual-programming technique in the Logic Editor, you link objects to create a fully working simulation, including dynamic elements with free-play capability.

At any stage in the development process, you can run the application to see how it functions. Rapid's multimedia interface supports MCI drivers for seamless integration of animation, digital audio and video, overlays, and scanners. A communications module lets you link to external programs and devices.
Price: US$6000

Contact: Emultek Ltd., Jerusalem, Israel, +972 2 870770;
fax: +972 2 870773.
Circle 1041 on Inquiry Card.

ARCserve supports popular tape and optical drives, including jukeboxes and other storage devices. ARCserve can back up all versions of NetWare, as well as connected DOS, Windows, Windows NT, OS/2, Mac, and Unix workstations and servers.

Price: call for prices
Phone: +44 1737 775500; fax: +44 1737 775520
Circle 1056 on Inquiry Card.

MICROSOFT MAIL UTILITY FOR BUSY PEOPLE

If you use Microsoft Mail and run full-screen applications, you know there is no simple way to find out when you have mail waiting until you complete the task at hand. WinSoft International (Marlow, U.K.) says that its Notifier utility rectifies this problem. The utility flags the arrival of incoming mail by overlaying a message window, through which you can opt to acknowledge receipt or elect to read your mail immediately. You can also configure

the utility to trigger an audible alarm, providing you have a suitable sound card and speaker driver.

Price: Starter kit, 10-user pack, £240; additional users, £12 each; unlimited site license, £1920
Phone: +44 1628 488444; fax: +44 1628 478560.
Circle 1424 on Inquiry Card.

MOOD MUSIC FROM YOUR COMPUTER ▲

You can play ever-changing and free-flowing harmonious background music from your PC with Koan Plus. The Koan Music Engine composes music in real time, offering a music kaleidoscope; pieces can even include preset seed melodies. Koan Plus comes with 10 pieces for you to play, 10 MIDI file recordings of these pieces, and 10 studio-mastered

mood music tracks.

The program from Koch Media (Farnham, U.K.) lets you record Koan piece outputs to a MIDI file. Koan Plus includes a number of modules, so you can easily load your pieces, put them in albums of your choice, and play them. With an analyzer module, you can see the way the music changes—bars move in time with the music, reflecting the changes as they happen. In addition, you can see the MIDI line usage or the range of notes Koan Plus uses.

Price: £29.78
Phone: +44 1252 714340; fax: +44 1252 711121.
Circle 1047 on Inquiry Card.

No

speed

limits

beyond

this

point...

DIGITAL TERRAIN MODELING UNDER AutoCAD

Topography and Digital Terrain Modeling, which runs under AutoCAD 11 and 12, assists you in all the steps of a topographical project (e.g., automatic drawing of contours and 3-D terrain vision). The program works with points imported from ASCII files in different formats and provides functions for inserting, deleting, interpolating, editing, filtering, and converting.

You can draw maps using point codes and define breakpoint lines using different criteria. The program provides graphic and analytical definition of alignments; calculation and representation of longitudinal and transversal sections, obtained from DTM or 3-D digitized maps; comparison of profiles in different states of evolution; and instant drawing of sections. You can perform volume calculations from meshes, DTM, or cross sections.

Price: US$990
Contact: TCP Informática y Topografía, Malaga, Spain, +34 5 2439771; fax: +34 5 2431371.
Circle 1043 on Inquiry Card.

What's New International Software

ALGORITHM FOR NEUDESK ▲

The Kohonen algorithm for Neural Computer Sciences' (Totton, U.K.) NeuDesk Windows-based neural network extends the package's capabilities for analyzing complex data sets. The algorithm implements Kohonen's concept for competitive learning, which lets a network discover patterns and associations in data without supervision. Applications for the Kohonen algorithm include data reduction, image and spectral analysis, and real-time alarm monitoring.

The Kohonen algorithm makes nodes in the grid compete for input patterns, modifying their weights during a training session to represent single data items or groups. In this way, Kohonen-based neural networks can cluster information and detect novelty, even when analyzing apparently unconnected data. You can produce a topological map of the data.
Price: US$50; with NeuDesk, US$385
Phone: +44 1703 667775; fax: +44 1703 663730.
Circle 1054 on Inquiry Card.

DOCUMENT IMAGING

Kofax Image Products' Ascent family of software modules consists of document-imaging components that allow systems integrators and VARs of midrange solutions to incorporate advanced imaging features. Available from Headway Technology Group (Aldershot, U.K.), the software includes Ascent Capture and Ascent ViewMana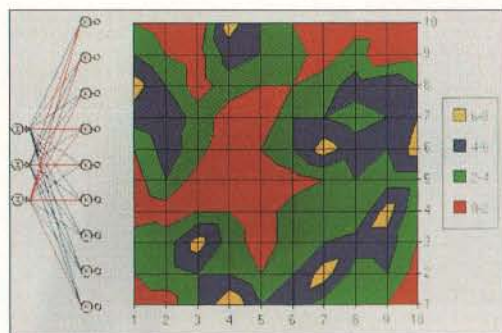ger. Ascent Capture can process 20,000 pages per day and gives mainframe levels of reliability and error recovery. It streamlines the entire process through features such as rated scanning and indexing aids.

You can integrate Ascent ViewManager into work-flow systems, imaging databases, and document management systems. It works with a database query front end that locates documents in storage and displays them by activating ViewManager. It's compatible with SQL databases.

Prefetching lets the system anticipate image requirements and automatically move successive pages of a long document from optical storage to local memory. Users can view image and non-image files simultaneously without leaving the application.
Price: Ascent Capture, US$2869; Ascent ViewManager, US$374
Phone: +44 1252 333575; fax: +44 1252 314445.
Circle 1059 on Inquiry Card

What's New

CONTROL AND DATA ▲ ACQUISITION SOFTWARE

PCVue2 2.13 includes a window editor to customize the user interface, user-configurable alarms, trend capabilities for logging data, report features for producing hard-copy data, and a communications manager for transferring data over a network. Arc Informatique's (Sèvres, France) object-oriented multiplatform software for supervisory control and data acquisition features on-line configuration capabilities that display modifications while the system supervises the process.

The program's open architecture uses DDE and NetDDE to access other programs (e.g., Microsoft Excel, Word, or Access) and DDLs to seamlessly integrate external software. PCVue2 2.13 can operate in two languages simultaneously and is available in a development version that lets you construct and change applications on-line without shutting down the system.
Price: from US$1000
Phone: +33 1 46230035; fax: +33 1 46238602.
Circle 1049 on Inquiry Card.

CRYSTAL REPORTS 4 ▶

Contemporary Software (St. Albans, U.K.) has added the following to Crystal Reports 4: faster report generation, 20 extra direct calls for the Report Engine, 40 VBX properties, OLE 2.0 compatibility, enhanced analysis facilities, and easier report design. Version 4 supports stored procedures, so you can

sort data on the SQL server. Enhanced analysis facilities give you the tools to turn mounds of low-level data into useful management reports. These include 12 styles of integrated graphics and a Select Records Interface that simplifies refining a query. Crystal Reports Experts prompt you for choices and automatically generate the report. Other design features include a Tab Interface, Best Fit button, Ruler, Automatic Styles, and new Grid control options.
Price: standard edition, £149; professional edition, £299
Phone: +44 1727 811999; fax: +44 1727 848991.
Circle 1426 on Inquiry Card.

3-D MODELER FOR MAC APPLICATIONS

The Amapi interactive graphics modeler for generating complex 3-D geometry combines intuitive modeling, accuracy, and power, according to Yonowat (Floirac, France). You can generate models by extruding basic 2-D lines, including NURBs and composite curves, and by using Boolean operations in combining 3-D primitives such as boxes, spheres, and parallelepipeds. To assemble models, you use rotation, copy, mirror, and other manipulation operations; to remove discontinuities and sharp edges, you use filleting, chamfering, and surface smoothing. A flattening tool creates and links a flat pattern to a 3-D object. Amapi stores your models and shapes in a graphics catalog, so you can quickly re-

trieve them for later use.
Price: FF 3950
Phone: +33 56869108; fax: +33 56866117.
Circle 1423 on Inquiry Card.

CLIENT/SERVER TESTING TOOLS ▶

A multiuser system and performance testing tool for client/server applications, LoadRunner/XL can test server functionality and performance under intense load conditions. The program simulates large numbers of clients—PC or Unix—interacting with the server by running thousands of user transactions in parallel. You can use the program for capacity planning and for tuning server performance under load by com-

paring results of several tests.
The program from Mercury Interactive (Gillingham, U.K.) can synchronize multiple user actions to test for data integrity. You can monitor and control an entire run, zoom in on an individual user, and load, run, stop, or mon-

itor any virtual user in the test session.
Price: LoadRunner Controller, £19,500; Virtual users, 50 (£13,250) to 3000 (£72,500)
Phone: +44 1634 262525; fax: +44 1634 262828.
Circle 1425 on Inquiry Card.

The Hardlock software protection system.
Universal and individual.

Across the world, there are currently over 1.5 million Hardlocks in use. They are all virtually hackproof, daisy-chainable, individually interchangeable, compatible, and with an almost incredible variety of forms and features, they are fast and easy to implement.

Want to know more? Want to find out which Hardlock will suit your application? Want to order the FREE Hardlock evaluation pack? No problem! Just dial...

MICROSOFT WINDOWS COMPATIBLE

yes It runs with NetWare

Available and Ready for OS/2

CE 93

COMPUTERIZED BOX-OFFICE SYSTEM ▶

Theaters, cinemas, and arts centers can computerize box offices with Asgard Ticket-Base. You can use the Windows program for arenas with up to 100,000 seats. In addition to providing standard box-office features, such as ticket printing and receipts reports, TicketBase offers contact management features for patrons and performers with automatic reminders, mail merge, and label printing. An event-planning module includes a visual calendar showing dates booked and free.

The package from Asgard Software (Oxford, U.K.) automatically calculates cost of tickets, as well as mail and credit charges; automatically checks for double bookings; lists all events by date, performer, location, or event name; and keeps full financial details.

You can create a seating plan by simply drawing it on-screen, design up to 10,000 locations with up to 100,000 seats each, and flag seats with special access. TicketBase also provides automatic reminders for any task connected to a patron or performer and financial and marketing reports.

Price: £495
Phone: +44 1865 209084;
fax: +44 1865 771461.
Circle 1427 on Inquiry Card.

MAXIMIZE YOUR MULTIMEDIA POTENTIAL

With MPC Wizard 3.0, you can make the most of multimedia on your PC, according to SoftKey International (London, U.K.). MPC Wizard 3.0 automatically checks the configuration of your multimedia system. Diagnostic tests check graphics, sound, video, CD-ROM drive, and Video for Windows performance. The program displays results on-screen. You can also save them to a disk or the Clipboard or print them. Some troubleshooting tips help you solve compatibility problems.

Price: £14.99
Phone: +44 181 789 2000;
fax: +44 181 789 5626.
Circle 1055 on Inquiry Card

SPARC COMPILERS

Associated Computer Experts' (Amsterdam, The Netherlands) Expert SPARC Compilers feature support for a variety of programming languages, a global optimizer, and a SPARC code generator. The global optimizer provides Markov-based execution frequency analysis, global register allocation, constant propagation, copy propagation, constant folding, strength reduction, and support for asynchronous control flow.

The SPARC code generator provides a peephole optimizer, TotalView and dbx support, and interfaces to Solaris as ld and libs. For testing and QA/QC, Expert SPARC Compilers provide the following: ACE/SCO SuperTest C Validation Suite, FSTC FORTRAN 77 Validation Suite, BSI Pascal Validation Suite, BSI Modula-2 Validation Suite, CADESE CASE, and ACE Benchmark Suite.

Price: NLG 899
Phone: +31 20 6646416;
fax: +31 20 6750389.
Circle 1422 on Inquiry Card

(038C/2)

With BW-MultiConnect and BW-Connect NFS for Windows NT You Won't Be Left Out In The Cold.

"Fred, now here's an on-line client/server solution that will warm up any application!"

It's easy to get your NFS clients and NetWare® users to warm up to the benefits of Windows NT.™ Reel them in with BW-MultiConnect™ and BW-Connect™ for Windows NT. Both solutions are from Beame & Whiteside of course, the experts in on-line Windows NT connectivity.

Windows NT users will bite on a wide range of NFS file and print services with BW-Connect NFS Client for Windows NT. It's the first multi-threaded NFS client for Windows. And it's kernel-mode,

delivering throughput up to 1,000 Kbps on some platforms. And BW-Connect NFS Server lets you share services with other NFS clients.

For all of your NetWare users, BW-MultiConnect for Windows NT is the catch of the day. With full NetWare server emulation, it provides open access to Windows NT file and print services. And best of all, you won't need to add any extra software on the NetWare client side.

Put your connectivity problems on ice with these affordable

Windows NT solutions. As Fred says, "If the software isn't from Beame & Whiteside, throw it back."

For your FREE 30-day personal evaluation just call or fax your distributor listed alongside this ad today.

Let's Connect!℠

Beame & Whiteside Software™

See us at NetWorld & Interop
Booth #2610, South Hall A.

Circle 529 on Inquiry Card.

yes
It runs with
NetWare.

BEST OF
BYTE
COMDEX Fall '93
Finalist
New NFS Server
for DOS & Windows

Microsoft®
SOLUTION PROVIDER
Partner

NETWORLD INTEROP
InteropNet
CONTRIBUTOR
ATLANTA
GOLD

Data Comm
MAGAZINE
TESTER'S
CHOICE
Best Advanced
Package: Windows

$E=MC^2$

One Great Solution...

Many Great Solutions...
and the Best Value Per Port in the Industry!

Discover Boca Research's amazingly affordable line of networking products that combine excellence in design and engineering with all of the latest features that make installing and using them faster and easier than ever. • Intel and Microsoft plug and play compatibility • High performance software drivers • Link status and programmable diagnostic LEDs • NE2100 or NE2000 compatible • Novell tested and approved.

From PCMCIA, PCI-, VL-, and ISA-Bus ethernet adapters, to 8-, 16-, and 24-port hubs, we deliver the no-compromise networking solutions that you and your customers demand. Each product is backed by a powerful five-year warranty and sold through a world-wide network of select marketing partners that includes leading distributors and resellers.

Our aggressive product support includes 24-hour bulletin board services through Boca Research and CompuServe. Convenient QuickFax provides fax-on-demand product information whenever you need it, and our technical support hotlines provide multiple choices in over-the-phone services.

From video graphic adapters to fax, data, voice, and multimedia communications products, to input/output devices and networking products, Boca Research leads the industry in product quality and responsive service. All this and more are part of our commitment to you.

FIVE YEAR
MADE IN USA
LIMITED WARRANTY

yes
NetWare
Tested and
Approved

Call your authorized Boca Research distributor today!

BOCA
RESEARCH INC

(407)997–6227 • Fax (407) 994–5848
1377 Clint Moore Road, Boca Raton, FL 33487-2722

Circle 501 on Inquiry Card (RESELLERS: 502).

What's New

FILE AND DATA MANAGER FOR THE AMIGA

Optonica (Lutterworth, U.K.) offers two packages for the Amiga: Info-Nexus, which is a file and data management system, and Data-Nexus, which is a high-specification multimedia database.

The Info-Nexus system features split-screen source/destination directories, copy, move, rename, clone, delete, file interrogation, file and embedded text searches, filtering, assignments, print controller, and media handling of virtually all Amiga file types. The program includes sound, CDXL video, and animation files direct from disk spooling; document editing with search; and disk relabel and format.

Data-Nexus features the following: multiple simultaneous database access; VCR-style controls; find and filter controls; visual print layout; mail merge and label printing; an on-line calculator; and media fields that support text, pictures, animations, sound, music, and CDXL video.

Price: £29.95 each
Phone: +44 1455 558282; fax: +44 1455 559386.
Circle 1052 on Inquiry Card.

SHAREWARE DATABASES

Ark Angles (Hazelbrook, Australia) offers shareware programs for Windows and DOS. FileIt for Windows 1.10 and FileIt for DOS 2.10 are free-form databases that can store records of just about anything you wish. The search facility allows you to find a string anywhere in your records, and the programs support records of up to 32 KB.

FileIt-GR for Windows 1.10 is a graphics database for storing, viewing, and sorting files. It supports search, import, export, and record titles of up to 21 characters. WriteIt 2.10 is a word processor and contact database that features drop-down 3-D menus, full mouse support, and quick keys on view at all times. The program includes word wrap and on-line help. K-Tools Utilities 1.1 lets you make or remove whole directory trees, copy files safely, synchronize whole drive or directory trees, encrypt and decrypt files, wipe files totally, and load any key into the keyboard buffer.

Price: call for prices
Phone: +61 47 588100; fax: +61 47 588638.
Circle 1421 on Inquiry Card.

KNOWLEDGE-BASED SYSTEM

Attar Software (Leigh, U.K.) has added new toolbars, drag-and-drop editing, and an enhanced object dictionary to XpertRule 3.0, its knowledge-based system. Additional features include the map, decision trees, pattern rules, optimization, and Xpert-Run.

A new toolbar of icons provides you with shortcut functions from the menus. You can drag attributes from dictionaries and drop them into tasks, resize columns by dragging margins, and move attributes and values around.

The dictionary holds attributes, variables, arrays, dialog boxes, and procedures as objects, which you can freely name and optionally tie to other objects. You can expand and contract the display of subtrees and repeated subtrees; create tasks with numeric, date, and Boolean outcomes; and set the inference mode to Decision tree or pattern rules. XpertRule 3.0 lets you exclude values from user selection, resize the tree display, and add attributes on the fly when building trees.

Price: US$395
Phone: +44 1942 608844; fax: +44 1942 601991.
Circle 1053 on Inquiry Card.

What's New International Software


```
                          WINlink 3
File  Disk  Tree  View  Psion  Tools  Options  Help

[C-] C:\WINDOWS\*.*                    [M-] M:\*.*
..           <DIR>                      AGN         <DIR>        04/07/93 10:10
LIGHTWIN     <DIR>      03/11/94 12:37  APD         <DIR>        12/12/94 16:17
MSAPPS       <DIR>      03/11/94 12:50  APP         <DIR>        12/12/94 16:04
SYSTEM       <DIR>      02/11/94 18:13  CVT         <DIR>        18/01/95 19:55
_DEFAULT.PIF   545  02/11/94 18:27      DAT         <DIR>        04/07/93 10:10
~SPL046C.TMP    86  29/11/94 16:47      FAX         <DIR>        18/01/95 19:55
~SPL0F72.TMP   300  29/11/94 16:44      GOBBLE      <DIR>        12/12/94 15:51
~SPL1718.TMP    86  29/11/94 16:45      GPK         <DIR>        28/12/94 18:08
~SPL301C.TMP 450346  29/11/94 16:45     KALL        <DIR>        05/01/95 16:09
~SPL3C1B.TMP   300  29/11/94 16:46      OPD         <DIR>        12/12/94 23:21
~SPL3D23.TMP 456456  29/11/94 16:47     PIC         <DIR>        12/12/94 16:20
~SRC016C.TMP     0  12/01/95 16:57      SPR         <DIR>        04/07/93 10:10
~SRC0A14.TMP     0  25/11/94 11:16      WDR         <DIR>        13/12/94 10:52
~SRC0C00.TMP     0  07/12/94 14:26      WLD         <DIR>        04/07/93 10:10
3BASE    .ANM   217  13/12/94 12:25     WRD         <DIR>        04/07/93 10:10
3BASE    .GRP  2576  24/01/95 10:18     WVZ         <DIR>        12/12/94 17:46
ACCESSOR.GRP 14903  24/01/95 01:18      SYS$TYCL.PDD  768  13/10/94 01:14

307 file(s), 11414702 Bytes              2 file(s), 3888 Bytes
Press F1 for Help.                                              NUM
```

DATA TRANSFER FOR PSION SERIES 3 AND 3A ▲

Widget Software (Knebworth, U.K.) has added features to Winlink 2, its data transfer utility for the Psion Series 3 and 3A. They include drag and drop, backup and restore, rescue of deleted records, automatic on-screen button help, and WordPerfect support. You can copy or move files with a single mouse movement. You can back up and restore whole SSDs or the internal disk M to your PC's hard drive. If you accidentally delete an address or telephone number from your database file, you can transfer the whole file to Winlink 2, which will restore all records that you have deleted since the last time you compressed your files.
Price: £29.99
Phone: +44 1438 815444; fax: +44 1438 815222.
Circle 1045 on Inquiry Card.

MOBILE E-MAIL SYSTEM

Remote WinMail 2.0 connects users to MHS-based E-mail systems, CompuServe, the IBM Mail Exchange Service, the BT Mailbox Service, the Dialcom messaging community, and the Concert Mail service, with worldwide access to public X.400 systems and fax terminals. The package lets you create and read messages off-line.

Remote WinMail 2.0's support for multiple communications protocols means that traveling users can communicate with their E-mail systems through one user interface. Available from Opus One (London, U.K.), Remote WinMail 2.0 provides true multilocation support.

You can collect your E-mail when and where you choose. Remote WinMail 2.0 features message batching for transfer during a single connection, dial-up local phone numbers for access to public networks by travelers, and support for high-speed modems.
Price: £199
Phone: +44 171 321 0402; fax: +44 171 839 8488.
Circle 1051 on Inquiry Card.

Blasts from the Past

DENNIS BARKER

5 Years Ago in BYTE

Big month for sizzling systems. Apple's Mac IIfx was our cover model; buzzing along with a 40-MHz 68030, it was the fastest Mac yet. We took our first look at IBM's new RISC workstations, the RS/6000 line, built around a new 32-bit

A couple of IBM RS/6000s

superscalar CPU. This replacement for the "lackluster" RT peaked at 40 MIPS, or about 2.5 times faster than the SparcStation 1.

Who would have thought then that four years later Apple would be the world's leading producer of RISC systems and that Apple and IBM would be allies? Apple and IBM! What's next—Sinatra and Bono?

The 486 was the hot new thing, but many folks were cranking up performance by swapping in 386 motherboards. We reviewed 23 of them. You could buy a 25-MHz caching board for $1100.

Our newly named State of the Art section looked at applications architectures, including two "blueprints for the '90s": IBM's Systems Application Architecture and Digital's Network Applications Support.

Our reviewer had a tough time tackling HP's New-Wave enhancement layer for Windows 2.11. He liked the iconic, object-oriented interface, the ability to share data between programs, and the Agent task automater. But the programming required to turn Windows applications into NewWave applications, the cryptic error messages, and snafu after snafu made it seem less than worthwhile.

10 Years Ago in BYTE

The big AI issue. Some of the top names in the field wrote about transplanting expert knowledge to a computer, massive parallelism as the key to simulating intelligence, a program that teaches Lisp programming, theories of learning, and that miracle that machines have yet to match: vision.

> "**When we** meet those aliens in outer space, will we be able to converse? . . . Yes . . . because we'll both think in similar ways. . . . All intelligent problem solvers are subject to the same ultimate constraints— limitations on space, time, and materials."
>
> **— Marvin Minsky**

Reverse engineering the brain. John Stevens, a professor of biomedical engineering, wrote about the possibility of digital hardware that copies the circuitry of the brain. While computers might be smart enough to ignore the O.J. trial, according to the professor's estimate "it would take a minimum of 100 years of Cray time" to simulate the processing power of the human eye.

Hey Dad, Did CP/M Machines Have, Like, Big Fins? "The micro industry is moving toward new standards of adequacy that will eventually relegate 8-bit CP/M computers to low-cost entry-level systems—or relics, like my neighbor's 1957 Studebaker."
— Pournelle's column

Our correspondent in Japan visited Hitachi to see the new S-810 supercomputer. The hardware deserved raves for its 600-megaflop performance, faster than a Cray. But the researchers admitted that without hand-tweaking their Fortran code ran no faster on the supercomputer than on their M-280 mainframe.

15 Years Ago in BYTE

On the cover: Hewlett-Packard's new barcode-reading wand. Editorial director Carl Helmers was a proponent of distributing software on paper, in barcode form.

Jef Raskin, info-appliance proponent and mind behind the Macintosh concept, wrote about the technical difficulties of turning a computer into a musician's amanuensis: It listens to what you play, then writes it down as notation.

In the News "Xerox has announced a new concept of processor-to-processor communications intended for an office environment. This novel concept is called 'Ethernet.'" Rumor has it that Xerox and Intel are developing an Ethernet chip.

Steve Ciarcia explained how to build a multiuser system out of five integrated circuits, one of which was the new 16-bit 8088. The postcard-size board held enough memory and peripheral interfaces to support two 300-bps terminals, each running a Tiny BASIC interpreter.

Jon Udell dissected **"Microsoft's long-awaited 32-bit OS/2 2.0."** His critical assessment of the not-yet-ready operating system as it was five years ago:

Cool
Flat memory model
Can run 32-bit applications on a server
Multitasking DOS and Windows sessions
MVDMs (multiple virtual DOS machines)
Threads support multitasking between and within applications

Not Cool
Can't run Windows binaries
Few native applications
"Spotty" device-driver support
Can't run DOS-extended programs under MVDM
Development tools lacking

Diagnosis: He liked it, and he wanted it, even though it was "still wobbly." And "a final release is many months away." How many months? It would be nearly two years and several big changes before **IBM** would ship 2.0.

Electronic musician Hal Chamberlin discussed advanced techniques for real-time music synthesis. With new, low-cost micros like Texas Instruments' 99/4 having built-in three-voice synthesizers, music applications of PCs "may soon approach the popularity of accounting, word processing, and games."

The New Texas Instruments TravelMate™ 5000

• 75 MHz Pentium processor • Two Lithium ion battery packs • Serial infrared port • 10.4" Active Matrix or 10.5" Dual Scan displays • 524 million bytes (=500MB) or 810 million bytes (=722MB) Hard Disk Drives • 8MB RAM, expandable to 32MB • Multimedia package: Built-in 16-bit sound, internal speaker and dual mode microphone*

Introducing
The First Notebook
To Maximize The Pentium's
Full Potential.

When the rush was on to introduce a notebook with a Pentium™ processor, Texas Instruments decided to do what others thought couldn't be done.

We created a notebook that maximizes Pentium performance by integrating *full PCI bus architecture in our TravelMate 5000.*

And for flexible connectivity, we designed a way to allow external access from the PCI bus to the latest peripherals.

It was a challenge we addressed for two simple reasons: to give users true desktop Pentium performance for *faster running software and expansion capabilities* for investment protection.

So now you have a notebook with smoother full-motion video and enhanced 3-D graphics. In addition, we designed the TravelMate 5000 to take advantage of the "plug and play" capabilities of Windows 95® when it becomes available.

But our engineers didn't stop there. *We added a second lithium battery* without sacrificing size, weight or eliminating a floppy drive.

In addition, wireless communication with other notebooks and desktops is very quick and easy with our integrated infrared capabilities.

The Texas Instruments 5000. For more information on what others thought you couldn't do with a notebook, simply call 1-800-TI-TEXAS (e-mail: 2ti@msg.ti.com).

EXTENDING YOUR REACH™

TEXAS INSTRUMENTS

Is There a God?

RICK COOK

Get a tight hold on your credulity. Physicist Frank Tipler is about to take you on Mr. Toad's Wild Ride through modern physics in search of the answer to the immortal question: Is there a God?

I won't keep you in suspense. Tipler's answer, rooted in current notions in quantum cosmology, is: "Not yet, but there will be." What's more, we will all be reincarnated and live forever in a computer simulation in the final fraction of a second before the universe recollapses in the big crunch. While you may not agree with Tipler's ultimate conclusions, you'll learn a good deal along the way and—if you approach it in the right spirit—have a lot of fun getting there.

When traveling with Tipler, the journey often looms larger than the destination. Along the way, we encounter the Bekenstein bound, which is the limitation on the number of quantum states in a bounded region, or to put it another way, the upper limit on information density, assuming each quantum state encodes 1 bit. We also meet the Jesuit paleontologist Pierre Teilhard de Chardin, who coined the term omega point and whose notion of "radial energy" bears a striking resemblance to the modern notion of information. We are treated to historical discourses on the ideas of eternal return and the immortality of the soul, and discussions of more modern concepts, such as the halting problem, Taub universes, and Cauchy spacelike hypersurfaces.

At the very least, Tipler deserves an "A" for originality. This is not another of those dreary tomes about how the new physics proves that Hindu theologians and Taoist alchemists had it right all along. Tipler is a serious and well-respected scientist, a professor of mathematical physics at Tulane University, with a firm grasp of the standard model of particle physics. True, he is speculating on a cosmic scale, but his speculations are rooted in the tools and concepts of modern science. Nor is this book dreary in any sense. Again and again, Tipler manages to surprise and delight you with novel approaches and new information.

The Physics of Immortality is an outgrowth of an earlier book Tipler wrote with John D. Barrow, *The Anthropic Cosmological Principle*, which argued that we exist because the universe is exactly the way it is. They pointed out in considerable detail that if the universe were almost any other way, life as we know it would be physically impossible. There have to be three large-scale dimensions of space, for example, for chemical compounds to fold themselves in the ways they must for life to exist. It is only a short step from the anthropic cosmological principle to the old argument from design for the existence of God—that God must exist because the universe was designed exactly for us. In the earlier book, Tipler and Barrow pointed this out only to denigrate it. Here, Tipler isn't so sure.

Tipler's God at the end of the universe is no scowling Yahweh or thunderbolt-slinging Jupiter, or even kindly, gentle Gaea. His Holy Spirit is the universal wave function of quantum mechanics, and his God is the most super of all the supercomputers, an information-processing construct that literally encompasses the whole universe.

One novelty for what is ultimately a work of theology is that Tipler offers series of testable hypotheses based on his theory. This is probably the only work ever written that posits as an eschatological proof that "the mass of the top quark must be 185 ± 20 GeV, and the mass of the Higgs boson must be 220 ± 20 GeV."

It has been truly said that in physics the line between genius, wise guy, and utter loon is often hard to draw. I'm not

LANDSCAPING MADE EASIER

It's that time of year again. The seed catalogs have arrived, the snow and frost are slowly receding, and that means it's time to follow through on those plans to landscape your property. 3D Landscape, the newest member of the Books That Work series, is an easy-to-use guide for playing what-if games for redesigning your yard.

Because gardening and landscaping always make the top 10 list of ways in which Americans spend their leisure time, a host of software programs have tried to help do-it-yourself designers make informed landscaping decisions. Unfortunately, most of those earlier programs contained little more than a simple database of plant names, linked to shade, sun, soil, and temperature requirements. Not very inspirational or informative.

Because 3D Landscape is liberated from the nongraphical constraints of the earlier DOS programs, it presents information in a much richer manner. To get your creative juices flowing, the program contains beautiful template gardens for small city lots, huge formal English estates, and "informal" country spaces. By dragging and dropping icons of shrubs, trees, or other objects (from a database of 400 species) onto a preexisting template or new workspace, you quickly get a feel for what works and what doesn't.

Once you've designed your landscape, including placement, for example, of fences, paths, a pool, a patio, or even a barbecue, you have the option of viewing your work from a variety of dimensions. You can even see what plants and foliage will look like as they grow over the years. You also see how they change during the seasons.

Reality enters the picture when it comes time to see what your design costs and how to implement it. The program calculates the cost of each item in your landscape. A how-to section with text, animation, photographs, and sound teaches you how to construct projects. 3D Landscape eliminates almost all your excuses for not tackling that landscape project.

—Rich Friedman

sure which side of the line Tipler lands on with *The Physics of Immortality*. Perhaps he has a foot in all three camps—a notion no stranger, and ultimately no less plausible, than some of the physics that Tipler lays out for your delectation. ∎

Rick Cook writes novels from his home in Phoenix, Arizona. You can reach him on the Internet or BIX at rcook@bix.com.

A DEVELOPER'S GUIDE TO GRAPHICS FILE FORMATS

ENCYCLOPEDIA OF GRAPHICS FILE FORMATS
by James D. Murray
and William vanRyper
O'Reilly & Associates, ISBN 1-56592-058-9, $59.95

The *Encyclopedia of Graphics File Formats* attempts to lasso the most popular graphics file formats into a concise reference. The fact that it includes a CD-ROM filled with reference material is evidence of just how ambitious this task is. The result is a valuable tool for developers.

It provides just enough technical information to understand graphics file formats on platforms ranging from DOS to Unix. While it discusses some background and theory behind compression methods and design characteristics, it quickly cuts to the chase, outlining the essential header and body structures for each format. Complete specifications obtained from the creators are provided on the CD-ROM. Code samples demonstrate techniques for decoding and manipulating selected formats.

With the abundance of commercial libraries available in this ever-evolving age of multimedia, it's easy to mistakenly regard this book as obsolete. Common wisdom states that it doesn't make sense to spend an inordinate amount of time developing your own filter library when you can simply plug in someone else's and move onto more deserving tasks. However, it's never quite that easy. There's always that one missing format, that one buggy filter, and the risk of relying on an outside source to keep up with new specifications and format upgrades.

Even if you do use third-party libraries, this book's treatment of the Windows Device-Independent Bit Map's palette organization will give you insight into the operation of your selected library. This knowledge will let you implement or tweak a particular filter with your own preprocessing or postprocessing to achieve the desired result.

The *Encyclopedia of Graphics File Formats* contains dozens of insights and tips. It addresses some staple animation and video specifications (e.g., FLI, DVI, and MPEG) and supplies vendor addresses. This book is a must-have reference source for developers of image-based applications. ∎

—*Brendan Daunt and Matt Trask*

Brendan Daunt is the author of Image Commander, a database for storing, managing, and transforming graphical-image files. He is a senior software engineer at and Matt Trask is president of Communica (Cape Cod, MA). Trask can be reached on the Internet or BIX at matt.trask@bix.com.

When you paid extra for SCSI bus performance, this probably wasn't what you had in mind.

Aeronics' new 68 Position IDC 0.025 Pitch Ribbon Cable Terminator will give your SCSI bus a boost.

Why does my bus need a boost?
Peripheral devices such as hard drives, CD-ROM drives, scanners, and printers that are interfaced with your computer have connectors manufactured by different companies that are connected to cable made by who-knows-who. Differences in peripheral cabling and connectors create signal impedance mismatches which can cause signalling errors, bottlenecks, and data transmission errors.

Aeronics has the answers.
Aeronics, the originator of Forced Perfect Termination (FPT), has added a new interconnect product to our family of high performance terminators that will drastically improve your SCSI bus performance. Aeronics' new 68 Position IDC 0.025 Pitch Ribbon Cable Terminator easily clamps onto your ribbon cable with a standard tool and replaces the low performance passive terminators frequently furnished by peripheral manufacturers. For those active negation drivers that generate higher than normal pull-up voltages that tend to overdrive the SCSI bus, Aeronics has developed the ALT-2ANHC. This new electrical configuration sinks unusual amounts of current due to charged cable capacitance and is available in most of Aeronics' interconnect products.

How does Aeronics boost my bus?
Aeronics' FPT active terminators purposely mismatch impedance higher and lower than the impedance of the transmission line, "forcing" the SCSI bus to operate as if the impedance between the host and peripherals is always matched. Because our terminators provide higher noise immunity, results are seen as improved data integrity, allowing your system to transfer data farther and faster while improving overall SCSI reliability.

Terminate your SCSI bus performance problems once and for all.
Aeronics' advanced terminator technology is available in a broad line of connection and electrical configurations that meet or exceed ANSI SCSI 1, 2, and 3 standards. Our passive, active, FPT, and differential terminators are 100% electrically tested, burned in, and delivered on time. With a minimum MTBF of 87,600 hours (10 years continuous duty), our quality terminators are backed by a limited lifetime warranty.

Aeronics' high performance products leave passive and copycat terminators in the dust. If you want the same performance that high-end OEM SCSI systems have, choose the same terminator from Aeronics and give your bus a boost. Call or fax Aeronics today for the location of the reseller or distributor nearest you.

For off-the-shelf delivery, call TCC at 1-800-TEC CABL (1-800-832-2225).

AERONICS INC. ™

Turning concepts into reality.

12741 Research Blvd. • Suite 500 • Austin, TX 78759
Tel (512) 258-8040 • Fax (512) 258-8441

BYTE/Osborne Books are Available at the Following Locations

ALABAMA
Madison
Madison Books
and Computers
PH: 205-772-9250
FAX: 205-461-8076

ARIZONA
Tempe
Computer Library
PH: 602-820-0458
FAX: 602-491-7971

CALIFORNIA
Berkeley
Cody's Books
PH: 800-479-7744 (In CA)
800-995-1180 (Outside CA)

Cupertino
Computer Literacy
Bookshops
PH: 408-973-9955

Stacey's Professional
Bookstore
PH: 408-253-7521
FAX: 408-253-5861

Irvine
Irvine Sci-Tech Books
PH: 800-229-9514
FAX: 714-733-0122

Los Angeles
ASUCLA Students Store
PH: 800-456-0019
PH: 310-206-4041
FAX: 310-825-0382

Opamp Technical Books
PH: 800-468-4322
FAX: 213-464-0977

Palo Alto
Stacey's Professional
Bookstore
PH: 415-326-0681
FAX: 415-326-0693

Stanford Bookstore
PH: 800-673-2348
FAX: 415-322-2021

Sacramento
Tower Books
PH: 916-481-6600

San Diego
San Diego Technical Books
PH: 800-346-0071
FAX: 619-279-5088

San Francisco
Stacey's Professional
Bookstore
PH: 800-926-6511
E-mail: staceysbk@aol.com

San Jose
Computer Literacy Bookshops
PH: 408-435-1118
E-mail: info@clbooks.com

San Luis Obispo
Earthling Bookshop Inc.
PH: 805-543-7951

El Corral Bookstore
PH: 805-756-5315
FAX: 805-756-5351

Santa Barbara
Chaucers Bookstore
PH: 805-682-6787
FAX: 805-682-1129

Earthling Bookshop
PH: 805-965-0926
FAX: 805-962-6620

UCSB Bookstore
PH: 805-893-2082

Stanford
Stanford Bookstore
PH: 800-533-2670
FAX: 415-322-1936

Sunnyvale
Computer Literacy
Bookshops
PH: 408-730-9955

COLORADO
Boulder
BiblioTek
PH: 303-443-7037

Colorado Bookstore
PH: 303-442-5051
FAX: 303-444-6604

Colorado Springs
McKinzey-White
Booksellers
PH: 719-590-1700

Denver
Auraria Book Center
PH: 303-556-3230
FAX: 303-556-3736

BiblioTek
PH: 303-534-3460

Tattered Cover Bookstore
PH: 800-833-9327
FAX: 303-399-2279
E-mail:
books@tatteredcover.com

Englewood
Softpro Books
PH: 303-740-7751
FAX: 303-740-8152

Longmont
United Techbook Co.
PH: 303-651-3184
FAX: 303-651-3405

CONNECTICUT
New Haven
Yale Co-op
PH: 800-ELI-YALE
FAX: 203-772-3665

**DISTRICT
OF COLUMBIA**
Washington D.C.
Reiter's Scientific
& Professional Books
PH: 00-537-4314
FAX: 202-296-9103

GEORGIA
Atlanta
Oxford Bookstore
PH: 404-262-9975
FAX: 404-364-2729

HAWAII
Honolulu
Honolulu Book Shops
PH: 808-847-5551
FAX: 808-841-4674

University of Hawaii
Bookstore
PH: 808-956-4338
FAX: 808-956-4323

ILLINOIS
Naperville
Books & Bytes, Inc.
PH: 708-416-0102
FAX: 708-416-0375

INDIANA
Ft. Wayne
Majerek's Readers World
PH: 219-482-1531

MASSACHUSETTS
Boston
Charlesbank Bookshops
PH: 617-236-7442
FAX: 617-236-7418

Waterstone's Booksellers
PH: 617-859-7300
FAX: 617-437-0997

Burlington
Softpro
PH: 617-273-2919
FAX: 617-273-2499
E-mail:
books@softproeast.com

Cambridge
Harvard
Co-Operative Society
PH: 617-499-2000
FAX: 617-868-7038

Quantum Books
PH: 617-494-5042
FAX: 617-577-7282
E-mail:
quanbook@world.std.com

Wordsworth Books
PH: 617-498-0080
FAX: 617-354-4674

Newton Highlands
New England
Mobile Bookfair
PH: 617-527-5817
FAX: 617-527-0113

Worcester
Tatnuck
Bookseller/Databooks
PH: 800-642-6657
FAX: 508-756-9425
E-mail:
databooks@delphi.com

MICHIGAN
East Lansing
Grand River Books
PH: 517-333-3666

Flint
Young & Welshans
Book Sellers
PH: 810-732-0626
FAX: 810-732-2393

MINNESOTA
Minneapolis
Baxters Books
PH: 612-339-4922
PH: 800-626-1049
FAX: 612-339-6134

NEW HAMPSHIRE
Hanover
Dartmouth Bookstore
PH: 800-675-3616 (in NH)
800-624-8800 (outside NH)
FAX: 603-643-5170

NEW MEXICO
Albuquerque
Page One, Inc.
PH: 800-521-4122
FAX: 505-294-5576

NEW YORK
New York
Barnes & Noble
Fifth Ave.
PH: 212-807-0099

Computer Book Works
Warren St.
PH: 212-385-1616
FAX: 212-385-8193

McGraw-Hill Book Store
Ave. of the Americas
PH: 212-512-4100
FAX: 212-512-4105

OHIO
Cincinnati
University of
Cincinnati Bookstore
PH: 513-556-1800
FAX: 513-556-5555

OREGON
Portland
Powell's Technical Books
PH: 503-228-3906
FAX: 503-228-0505

PENNSYLVANIA
Pittsburgh
Book Center -
U. of Pittsburgh
PH: 412-648-2321
FAX: 412-648-1902

RHODE ISLAND
Providence
Brown Bookstore
PH: 401-863-3168
FAX: 401-863-2233

TEXAS
Dallas
Taylors Bookstore Ltd.
PH: 214-357-1700
FAX: 214-956-8012

VIRGINIA
Vienna
Computer Literacy
Bookshops
PH: 703-734-7771
E-mail: sales@tc.clbooks.com

WASHINGTON
Bellevue
Tower Books
PH: 206-451-1110
FAX: 206-454-0453

University Bookstore
PH: 206-646-3300
FAX: 206-634-0810

Seattle
University Bookstore
PH: 206-634-3400
FAX: 206-646-3340

WISCONSIN
Madison
University Bookstore
PH: 608-257-3784
PH: 800-993-2665
FAX: 608-257-9479

Milwaukee
Harry W. Schwartz
Bookshops
PH: 414-274-6460
PH: 800-236-7323
FAX: 414-274-6408

INTEL'S

Intel's successor to the Pentium is an impressive microprocessor, a bold new

generation of CISC. But is it enough to keep AMD, Cyrix, and NexGen at bay?

TOM R. HALFHILL

Fight with the big kid, and you'll get a bloody nose. That's the message Intel wants to send to rival chipmakers Nex-Gen, Cyrix, and AMD. Intel's telegram is a massive next-generation CPU soon to appear in high-end workstations and servers. The new chip has a name: P6.

The P6 contains the most transistors ever seen in a mass-market processor. It is aggressively superpipelined, supports out-of-order execution, and has a three-way super-scalar microarchitecture. These techniques are impressive but not original—new chips from NexGen and AMD sport similar designs. But Intel has a key advantage: The P6 has a unique secondary cache that is connected to the CPU by a dedicated bus. This 256-KB SRAM (static RAM) cache, which resides in the second compartment of the P6's unusual dual-cavity package, will dramatically simplify the design and construction of systems based on this chip.

Intel's estimated benchmarks peg the P6 at about 33 percent faster than a Pentium running at the same clock speed. (The first P6 will run at 133 MHz, which will place it about twice as fast as today's fastest 100-MHz Pentium.) The P6 is thus faster than NexGen's Nx586, which is shipping today, but in the same ballpark as AMD's K5 and Cyrix's M1—both of which should ship at about the same time as the P6.

Difficult Questions

The P6 is exciting news, but its unveiling also raises a number of serious questions. Is the P6 really a full generation beyond the Pentium? Are Intel's competitors finally catching up? Which CPU is the safest choice for reliability and compatibility? Which is the best buy in terms of price/performance? And looking further ahead, how do these latest CISC chips compare to their RISC counterparts? Will the x86 family remain competitive with such RISC chips as the PowerPC in Apple's Power Mac?

Some of these questions can't be answered until more of the competitive x86 chips hit the market. NexGen's Nx586 series has been shipping since last September, but the P6, K5, and M1 are not expected to be available in systems until the second half of this year. Until then, firm pricing and actual system performance are frustrating unknowns. But some questions can be answered today.

If Intel's competitors can meet their delivery schedules and performance promises, it appears they have caught up—at least temporarily. Measured against RISC chips, the P6 still isn't the fastest microprocessor you can buy, but it still runs the most software, and it continues to keep the performance gap narrow enough to deny the RISC machines an overriding advantage. The fastest general-purpose RISC processor is Digital Equipment's Alpha 21164, and its integer performance is only about 1.5 times better than the P6's.

Another CISC/RISC Hybrid

If it's true that great minds think alike, perhaps it's no surprise that the microarchitecture of the P6 resembles those of the Nx586, the K5, and, to a lesser extent, the M1. Engineering teams separated by thousands of miles independently designed each of these x86-compatible processors at roughly the same time. All four combine key elements of RISC technology with some novel concepts to define a bold new generation of CISC.

The P6 is obviously Intel's sixth-generation x86, the

The P6: What's New

- **5.5 million transistors in the CPU core, 15.5 million in the secondary cache**

- **Three-way issue superscalar microarchitecture**

- **14-stage superpipeline**

- **Five parallel execution units: two integer, one load, one store, and one FPU**

- **8-KB two-way set-associative primary instruction cache; 8-KB four-way set-associative primary data cache**

- **256-KB SRAM (static RAM) secondary cache on a dedicated full-speed bus, closely coupled to the CPU die in a 387-pin dual-cavity PGA ceramic package**

- **Transactional I/O bus and nonblocking cache hierarchy**

- **Out-of-order execution, dynamic branch prediction, and speculative execution**

- **Claimed performance: 200 SPECint92 at 133 MHz**

- **2.9-V, 0.6-micron, four-layer metal BiCMOS process technology**

- **CPU die: 306 mm square; SRAM die: 202 mm square**

- **Estimated power consumption: 20-W peak at 133 MHz**

P6

successor to the Pentium. Until now, the other chips have been considered fifth-generation rivals of the Pentium. But in terms of concept and performance, they are not very different from the P6. This is mostly because Intel's competitors are pursuing their own courses, no longer content to merely clone Intel's designs. By similar but independent routes, they are moving toward the same destination, and the once-discrete generations are starting to overlap.

All these processors are 32-bit implementations of the Intel x86 architecture, but their most interesting similarity is the degree to which they adopt the same techniques found in state-of-the-art RISC chips. This is largely because silicon-based microprocessors are a maturing technology, with the easiest gains in performance already behind us. Engineers at all the leading companies are resorting to exotic tricks to squeeze more performance from existing architectures, and even more exotic architectures are taking shape on the drawing boards (see the text box "The P7

and Beyond" on page 54).

In this respect, the P6 is a significant improvement over the Pentium. Although the Pentium broke ground by introducing the first superscalar x86 microarchitecture, it is a rather pedestrian superscalar design: a pair of integer pipelines that can process two simple instructions in parallel and in order, with no speculative execution. By contrast, the P6 is a three-way superscalar machine that can simultaneously dispatch five instructions internally. To amplify that wider instruction bandwidth, the P6 also has improved caching, more registers, better branch prediction, deep speculative execution, and out-of-order instruction handling. Clearly, it is a design leap beyond the Pentium.

It is not, however, a leap beyond its contemporaries. In some cases, the P6 is somewhat more sophisticated than the Nx586, K5, and M1; in other cases, somewhat less. All things considered, the P6's design is on a par with those of the other CPUs, with

P6 Technology

5.5 million transistors in the CPU core, 15.5 million in the secondary cache

the K5 its closest relative.

Another striking similarity is how the P6, K5, and Nx586 take the same approach to CISC/RISC hybridization. The P6 welds a traditional CISC front end to a RISC-like core, bridging the gap with a clever decoder that breaks down the lengthy CISC instructions into simpler operations that more closely resemble RISC instructions. The simplified operations—Intel calls them *micro-ops*—are then fed into a core that takes advantage of the latest RISC innovations. Micro-ops are easier to dispatch and execute in parallel than their complex x86 antecedents.

NexGen, the newest contender among x86 makers, was actually the first to introduce this idea in an x86 processor. NexGen refers to the simplified operations as *RISC86* instructions, and the Nx586 can issue up to four of them simultaneously to its multiple execution units. AMD's K5 has a similar four-way decoder, but the company refers to the simplified instructions as *R-ops* (RISC operations).

Whether you call them micro-ops, RISC86 instructions, or R-ops, the goal of this fission is the same: to overcome the historical limitations of the x86 instruction set while maintaining compatibility with existing x86 software. (Although Cyrix's M1 does not handle x86 instructions in this way, it adopts other RISC techniques to achieve comparable efficiencies.) On the outside—to programmers and the code they write—these processors look like standard x86-compatible CISC chips. On the inside, however, they work much like a modern, streamlined RISC chip.

Intel's term for this CISC/RISC hybrid instruction flow is *dynamic execution*. You'll find the same basic mechanisms if you pry off the lids of the latest RISC processors, including the IBM/Motorola PowerPC 604, the PowerPC 620, the Sun UltraSparc, the Mips R10000, the Digital Alpha 21164, and the Hewlett-Packard PA-8000.

Two Chips in One Package

The P6's most fascinating feature is a closely coupled secondary (level 2, or L2) cache in the same package with the CPU. Yes, the P6 is two chips in one package. One die is the CPU core, including two 8-KB primary (level 1) caches; the other die is a 256-KB SRAM that functions as the four-way set-associative L2 cache.

These two dies share the same 387-pin ceramic package but otherwise are separate wire-bonded chips. Some other companies

How Much Faster Is the P6?

Intel says the P6 is about twice as fast as the Pentium. But Intel is comparing the estimated performance of a 133-MHz P6 (200 SPECint92) to that of a 100-MHz Pentium (112.7 SPECint92). Is that fair?

"What you have to compare is given the same process technology, how one microarchitecture technology compares to the other one," says Lew Paceley, P6 marketing director. "A 100-MHz Pentium processor compares to a 133-MHz P6. Why? Because the P6 is superpipelined, so its clock rate is higher. On any given process technology, the P6 will run faster."

The P6's pipeline has up to 14 stages, compared to the Pentium's five stages. Yet the P6 can move instructions through its pipeline faster because the stages have shorter latencies. Of course, the P6 has many other advantages as well, such as three-way pipelines, a closely coupled secondary cache, a transactional I/O bus, more execution units, and so on. These are exactly the kind of features that distinguish one microarchitecture from another.

Intel acknowledges it is possible to fabricate a 133-MHz Pentium on the same process technology used for the P6 (0.6-micron, four-layer metal BiCMOS). Last year, in fact, Intel presented a paper at the International Solid-State Circuits Conference that described a 150-MHz Pentium on this process. Whether such a chip ever becomes a commercial product is another question, but Intel is known to have a 120-MHz Pentium that may show up in products soon and is rumored to have a 133-MHz version.

In any scientific experiment, the goal is to control all the variables except the one you're testing. If you compare two microprocessors that implement the same general architecture on the same process technology and both chips are running the same benchmark program at the same clock speed, any difference in performance must be due to the relative efficiencies of their microarchitectures.

The estimated performance of a 133-MHz Pentium is 150 SPECint92. Therefore, a P6 running at the same clock speed is about one-third faster, not twice as fast. Nevertheless, Intel is correct in saying that on any given process technology it's possible to make a faster P6. So if Intel ever does produce a 133- or 150-MHz Pentium on its 0.6-micron process, the P6 could be clocked at about 200 MHz—and it would be about twice as fast as the 133- or 150-MHz Pentium.

Intel P6 Block Diagram

Main memory

Packaged L2 cache
256 KB

Memory order buffer

Frontside bus
(64-bit divided
synchronous)

Backside bus
(64-bit core
synchronous)

(36 bits)

Address bus

Bus interface unit

Data cache
8 KB

Memory
interface unit

Data-result buses (up to 86 bits wide)

Fetch two cache
lines per clock cycle,
32 bytes per line

Address-
generation unit

Instruction fetch unit

Instruction cache
8 KB

Address-
generation unit

Branch
target
buffer
(512 entries)

Reservation station

Integer unit

x86 instruction decoder

Integer unit

Simple
decoder
(1 micro-op
per clock cycle)

Simple
decoder
(1 micro-op
per clock cycle)

Complex
decoder
(1 to 4 micro-ops
per clock cycle)

Microcode
instruction
sequencer
(greater than
4 micro-ops)

FPU

Register alias table

Reorder buffer
(40 physical registers)

Retirement register files
(14 architectural registers)

The P6 fetches x86 instructions from the executing instruction stream or from a speculative address, based on predictions by the branch target buffer. A parallel decoder converts the instructions into RISC-like micro-ops and assigns them to registers in the reorder buffer. The reservation station dispatches the micro-ops in parallel and out-of-order to five execution units. After completion, results are committed to the retirement register file and restored to original program order.

refer to this type of package as an MCM (multichip module), but Intel is calling it a *dual-cavity PGA* (pin-grid array). The difference is subtle and probably has more to do with marketing than technology, because MCMs have a reputation for being expensive. (Until Intel announces pricing for the P6 later this year, we won't know if the new terminology will improve that reputation.) No matter what it's called, the P6 is the first microprocessor intended for large-scale production that includes two dies in a single package.

The transistor counts are impressive: The CPU die has about 5.5 million transistors, and the L2 die has about 15.5 million. By comparison, the latest Pentium cores have about 3.3 million transistors. Of course, that doesn't include an L2 cache, because as with almost all microprocessors, the Pentium requires an external set of memory chips to implement a secondary cache.

If you're wondering why the L2 die in the P6 requires so many transistors for only 256 KB of memory, remember it's

SRAM, not DRAM. DRAM chips use only one transistor per bit and are periodically refreshed, and SRAM chips use four or six transistors per bit so that they don't need refreshing. That's what makes SRAM faster and more expensive than DRAM. The SRAM in the P6 has six transistors per bit.

Although the L2 die in the P6 has about three times as many transistors as the CPU die, the CPU die is physically larger. It measures 306 millimeters square (compared to 295 for the first Pentium), while the L2 die is only 202 mm square. Both are fabricated on the same process (0.6-micron, four-layer metal BiCMOS), but it's easier to pack transistors more densely onto memory components than onto logic parts.

Here's why Intel incurs the extra expense of packaging the CPU core and secondary cache together: First, the combination package will make it easier for

P6 Technology

256-KB SRAM (static RAM) secondary cache on a dedicated full-speed bus, closely coupled to the CPU die in a 387-pin dual-cavity PGA ceramic package

vendors to design high-performance systems around the P6. One challenge of designing a computer with a fast processor is fine-tuning the size and configuration of the secondary cache to the requirements of the CPU. Unfortunately, the low margins of today's PC-clone market do not encourage lavish R&D budgets. At the same time, clone vendors are under great pressure to get their machines out the door as soon as possible. The P6's built-in secondary cache is already tuned to the CPU and will let systems designers quickly integrate the package onto motherboards.

Another important reason for the dual packaging is enhanced performance. The secondary cache is closely coupled to the CPU core by a dedicated bus that is 64 bits wide and runs at the same speed as the CPU clock. If the CPU core is clocked at 133 MHz (at which speed the P6

SUPER CONTRAST!

DESIGNED FOR PROFESSIONALS

When you decide to upgrade to a 17", 20" or 21" monitor, a ViewSonic Professional Series monitor, feature for feature and dollar for dollar, is the most intelligent choice you can make. Incorporating leading-edge technology, these high performance monitors include our unique **Super Contrast** screen. The feature-rich ViewSonic 21PS, 20PS and 17 monitors provide screen images with phenomenal color contrast and clarity.

Our **Super Contrast** screen uses advanced technology to create sharp, rich images, starting with a special glass not found in most monitors. This glass minimizes the diffusion of light between the individual picture elements on the screen, virtually eliminating color washout caused by pixel crossover, producing richer, purer colors *(see illustration)*. To guarantee that the image is brilliant and crystal clear, a specially designed shadow mask ensures that more energy reaches the phosphors, and more light reaches you. Does all this advanced technology matter? It does when you see the brilliant, true-to-life images it produces.

Be assured the feature-rich ViewSonic Professional Series monitors will optimize your productivity. For example, all three monitors provide resolutions up to 1,600 x 1,280. And, with maximum refresh rates of 160Hz, you know the image and color on the screen will be stable, consistent and picture-perfect.

Plus, these monitors are equipped with our OnView™ control system which allows you to adjust images via an on-screen menu.

ViewSonic's Super Contrast screen eliminates pixel crossover (top), providing distinct pixelization (bottom).

ViewSonic . . . always on the leading edge of technology.

ViewSonic®
See The Difference!™

Tel: (800) 888-8583 or (909) 869-7976; Ext. 217
Fax: (909) 869-7958   Call FaxSonic at (909) 869-7318 (fax-on-demand)
Request Doc. 162 (21PS), 156 (20PS) and 152 (17)
Applelink: VIEWSONIC   Compuserve: 73374, 514

All products and trademarks are brand names of their respective companies. Specifications subject to change without notice.

Circle 109 on Inquiry Card (RESELLERS: 110).

will debut), the secondary cache also runs at 133 MHz.

The first Pentiums were clocked at 60 and 66 MHz, and they addressed their secondary cache over a 64-bit bus that ran at the full speed of the CPU core. However, as Pentium clock speeds increased, it became too difficult and expensive for systems designers to match those speeds on their motherboards. So the faster Pentiums use a clock divider to run the external bus at a fractional speed. For example, a 100-MHz Pentium runs its external bus at 66 MHz, and a 90-MHz Pentium runs its external bus at 60 MHz. The Pentium also uses this same bus to access main memory and other peripherals, such as the PCI chip set.

Because the P6 has a dedicated bus for the secondary cache, it achieves two efficiencies: full synchronous bus speeds and no contentions with other I/O operations. The dedicated L2 bus, or *backside bus*, is completely separate from the external I/O bus, or *frontside bus*, so the P6's secondary cache doesn't have to share access with main memory and peripherals. The 64-bit frontside bus can run at one-half, one-third, or one-quarter the speed of the P6's core, depending on the wishes of the systems designer, and the backside bus continues to run independently at full speed.

These are big improvements over the bus designs of the Pentium and most other x86 processors. However, NexGen was there first with a similar design. Although the Nx586 doesn't package the secondary cache with the CPU, it does integrate an L2 cache controller and a full-speed backside bus. Like the P6, the Nx586 accesses main memory and peripherals concurrently over a separate I/O bus that runs at a clock-divided frequency.

Digital's Alpha 21164, an exotic RISC chip, goes a step further by integrating a 96-KB secondary cache with the CPU core, just like a primary cache. This provides unprecedented cache performance at the expense of considerable silicon real estate—the Alpha 21164 contains a whopping 9.3

Increased Performance

Since 1978, Intel's x86 processors have steadily increased in performance. Although Dhrystone MIPS are no longer considered a good measure of CPU performance, they are the only benchmarks that span all six x86 generations.

million transistors, most of which are arrayed in the caches. The payoff is that the Alpha 21164 delivers 330 SPECint92 at 300 MHz. Intel's P6 isn't quite that fast (Intel estimates 200 SPECint92 at 133 MHz), but it strikes a better balance between speed and cost for its intended market.

One caveat: The P6's unusual design will make it difficult to calculate its price/performance ratio, even when we know the price. The built-in secondary cache is bound to make the P6 look more expensive than its competitors, but most of the other processors require an external set of memory chips, as well as a cache controller. The more efficient design of the P6's cache structure means the other processors will probably need more than 256-KB cache to achieve comparable performance. Keep those factors in mind when evaluating systems.

Intel's unique packaging also leaves the door open for more design variations. Future versions of the P6 could easily enlarge the secondary cache or take the more conventional approach of separating the cache from the package. A P6 with an external secondary cache wouldn't be pin-compatible with a multichip version because it would require 72 more pins (64 for the backside bus and eight for error-correction coding), but it would be just as fast, especially when burst-mode

SRAMs become available. History shows that Intel likes to offer numerous variations of a basic processor design. It provides more flexibility for system vendors and leaves less room for competitors to carve out market niches.

One variation that wouldn't make sense, according to Intel, is to connect an external set of memory chips to the P6's frontside bus to implement a level 3 cache. The Alpha 21164 is designed to work with an L3 cache, but Intel says it wouldn't buy much performance for the P6 because the L1 and L2 caches already are closely matched to the demands of the core.

P6 Microarchitecture

Modern microprocessor design requires careful attention to balancing all the different elements in the chip. It's not just a matter of adding more resources in the form of larger caches, wider buses, faster clocks, multiple execution units, and additional registers, because some of those resources will be wasted if they exceed the ability of others to keep up. Intel started designing the P6 about four years ago (roughly halfway into the Pentium's design cycle), and the Oregon-based engineers spent much of that time fiddling with simulators to determine the optimum combination of I/O bandwidth and execution capacity.

The closely coupled secondary cache is only one example of how Intel boosted the P6's bandwidth to feed the hungrier appetite of its superscalar core. Another example is how Intel made the primary and secondary caches nonblocking without increasing their size.

Each cache in the P6 is 8 KB; the instruction cache is two-way set-associative, and the data cache is four-way set-associative. The Pentium also has a pair of 8-KB primary caches, although both are two-way set-associative. Normally, you would expect to find larger primary caches (and a larger secondary cache) in a three-way superscalar machine like the P6. Instead, Intel has designed the cache hierarchy to allow requests to flow

More Complex Designs

Since the 8086's debut in 1978, the number of transistors on Intel's x86 microprocessors has increased by a factor of about 190. This reflects the greater complexity of Intel's designs.

Give Me

Liberty

Chip by Chip, Feature by Feature

	INTEL P6	INTEL PENTIUM[1]	AMD K5	CYRIX M1	NEXGEN NX586
Clock speed	133 MHz	100 MHz	100 MHz	100 MHz	93 MHz
Performance (SPECint92)[2]	200[3]	112.7	147[3]	147 to 169[3]	112.7[3]
Superscalar design	Three-way	Two-way	Four-way	Two-way	Three-way[4]
Execution units	Five	Three	Five	Four	Three[4]
Branch prediction	Dynamic	Dynamic	Dynamic	Dynamic	Dynamic
Speculative execution	Yes	No	Yes	Yes	Yes
Out-of-order execution	Yes	No	Yes	Yes	Yes
Transistors	5.5 million[5]	3.3 million[6]	4.3 million	3.3 million	3.5 million
Level 1 cache	16-KB split	16-KB split	24-KB split	16-KB unified	32-KB split
Level 2 cache	256 KB packaged	External	External	External	External[7]
Process technology	0.6-micron BiCMOS	0.6-micron BiCMOS	0.5-micron CMOS	0.65-micron CMOS	0.5-micron CMOS
Die size	306 mm square	163 mm square	Unknown	394 mm square	196 mm square[8]
Availability	Second half of 1995	Mid-1994	Second half of 1995	Second half of 1995	Late 1994
Price (in lots of 1000)	Unknown	$673	Unknown	Unknown	$569

Notes:
[1] P54C Pentium at 100 MHz.
[2] Benchmark suite revised November 1994.
[3] Estimated performance.
[4] Not including optional FPU.
[5] Packaged level 2 cache has 15.5 million transistors.
[6] Pentium-60 and -66 have 3.1 million transistors.
[7] Integrated level 2 cache controller.
[8] NexGen plans die shrink to 119 mm square in April.

from one cache to the other without blocking execution.

If the P6 cannot find the instructions or data it needs in the primary caches, it quickly checks the secondary cache. Thanks to the backside bus, the penalty is only one clock cycle. If the P6 can't find what it needs there, it must access main memory over the frontside bus. The penalty for a double cache miss is usually considerable, because the frontside bus runs at a divided clock speed and main memory is relatively slow. But Intel has taken steps to minimize the impact.

First, as described above, the independent frontside and backside buses operate concurrently. Second, the frontside bus is *transactional*, another improvement over the Pentium.

While the P6 is waiting for a memory access to complete, it can begin another access. As many as eight of these transactions can be outstanding. Transactions can be loads or stores, and each transaction is recorded in a MOB (memory order buffer) so that the CPU won't mix them up. In fact, the P6 can let a load "pass" a pending store, which greatly accelerates throughput. While a Pentium would stall in this situation, the P6 can continue issuing loads and stores and execute other nondependent instructions that are queued in its buffers.

To maintain coherency between the caches and memory, the P6 implements the MESI (modified, exclusive, shared, invalid) coherency protocol. This guarantees cache coherency in both single- and multiprocessing systems, and it also guards against potential tragedies caused by self-modifying code. Many old x86 programs (and even some naughty new ones) use this technique, which can wreak havoc when the program rewrites portions of itself that are already cached. The P6 detects this, and even dedicates a machine-specific register to track how often it happens. An operating system could check this register and warn you if your programs are too clever for their own good.

The combination of nonblocking caches, separate fetch buses, and an eight-entry transactional buffer goes a long way toward ensuring that the P6's multiple execution units won't run out of work to do.

Instruction Fission

One of the historical burdens of the 17-year-old x86 architecture is its complex instruction set. In the late 1970s, when Intel engineers were designing the original 8086, the goal was to pack the richest possible instruction set onto a chip with only 29,000 transistors. They did it by creating complex instructions that do a lot and require microcode. But as processors grew more powerful, those long, complex instructions became a liability. RISC processors take the opposite tack by using instructions that are relatively simple and are hard-wired onto the chip.

RISC instructions are easier to handle internally, especially in a superscalar machine, so the trend in CISC design is to make the complex instructions appear more RISC-like without changing them. Changing the CISC instructions is unthinkable because it would render existing software obsolete.

The approach introduced by NexGen and adopted by AMD and Intel is to build an x86 preprocessor with a sophisticated decoder that parses the complex instructions into fragments of simpler operations: These are the RISC86 instructions, R-ops, or micro-ops mentioned before. The simplified operations are easier to dispatch and execute in parallel than full-length x86 instructions.

Instruction fission begins soon after the x86 instructions are fetched into the pipeline. Interestingly, the P6 adopts another RISC technique at this point that the latest RISC chips are retreating from: superpipelining. The P6 has a 14-stage pipeline (although in some instances stages can overlap), considerably longer than the Pentium's five-stage pipeline. (Due to the complex nature of the x86, some instructions consume additional pipeline stages, and floating-point instructions may require many more stages.)

Superpipelining is falling out of favor with some engineers because it can impose costlier penalties when the CPU misses the cache or mispredicts a branch. Intel is trading off those disadvantages against faster overall throughput, because the P6's stages have shorter latencies than the Pentium's, and they help to keep the execution units busy.

The first stage of this superpipeline calculates the next value of the instruction pointer, which

P6 Technology

2.9-V, 0.6-micron, four-layer metal BiCMOS process technology

determines the location from which the CPU will fetch the next instruction. This calculation isn't as simple as merely incrementing the pointer, of course, because there might be a pending interrupt or a branch that would move the pointer to a different location.

During the next three stages, the P6 fetches two cache lines of 32 bytes each, marks the boundaries of each instruction, and passes 16 aligned bytes to the decoder. Two cache lines are fetched per clock because x86 instructions can be long, so there's a good chance an instruction will wrap around from one cache line to another. Marking the boundaries is important because x86 instructions vary in length (8 to 120 bits), so it's not obvious where one instruction ends and the next begins. (RISC instructions are easier to fetch because they're always 32 bits long.)

The following stages of the P6's pipeline are noticeably different from the

> **P6 Technology**
>
> Five parallel execution units: two integer, one load, one store, and one FPU

Pentium's. Instead of merely decoding the x86 instructions and shoving them down a pair of execution pipelines, the P6 converts them into sequences of micro-ops and appends various flags and status bits to prepare them for superscalar out-of-order issue. To do this, the P6 has a decoder unit that is much more sophisticated than the Pentium's; it actually consists of three decoders working in parallel.

Two of these decoders parse the simpler x86 instructions, mostly integer operations. When they encounter a more complicated x86 instruction, they pass it along to a more sophisticated third decoder. Some instructions are too complex for even this decoder to parse, so they're referred to a microcode instruction sequencer, which corresponds to the microcode ROM in earlier x86 chips.

In any case, the final output is the same: triadic micro-ops that always contain three operands—two sources plus a destination.

They're longer than true RISC instructions because they have to accommodate the more complex x86 floating-point operations, but they're definitely easier to manage than the original x86 instructions from which they are derived.

Most x86 instructions map directly into single micro-ops. All but the most complex instructions translate into one to four micro-ops. When an instruction is so complex that it must be broken into more than four micro-ops, it is passed to the microcode instruction sequencer. The worst-case situation is difficult to pin down because x86 instructions can have prefixes that greatly multiply their complexity. However, if the instruction doesn't have a repeat prefix, isn't looped, and isn't a string operation, the worst case is that it will translate into 204 micro-ops.

The two simpler decoders in this mechanism can each generate one micro-op per clock cycle. The third decoder can generate up to four micro-ops per clock cycle. Therefore, with each clock cycle, the unit

The Shape of Systems to Come

ANDY REINHARDT

Computers built around the Intel P6 chip will initially resemble today's top-of-the-line Pentium boxes, but they will have an even more upscale configuration. Expect to see at least 1-GB hard drives, 32 MB of RAM or higher, and high-end graphics controllers for workstation applications. Some servers will be multiprocessors.

Over the next two years, however, P6 systems will depart from their Pentium predecessors. Desktops will possess more storage and multimedia horsepower than all but today's snazziest workstations. Better built-in diagnostics and error reporting in the P6 will make it easier for vendors to build reliable, enterprise-class servers. And the P6's improved SMP (symmetric multiprocessing) support, in conjunction with new SMP-capable versions of OS/2 and NetWare, will fuel a surge in high-end servers.

The Next Stage

Intel had expected servers to be the first P6 systems available, with desktop systems not far behind. However, it appears that both servers and desktop units will hit computer stores almost simultaneously. "We will definitely have P6 on the server first, but desktops will immediately follow," says Jeffrey Lu, principal architect for Digi-

tal Equipment's PC business unit premium desktop line.

Hewlett-Packard sees immediate demand for upscale P6 desktop models. Other companies jumping on the P6 bandwagon include ALR, AST, AT&T Global Information Systems (the former NCR), Dell, Gateway, and Unisys. The first desktop systems will start in the $4000 to $4500 range and climb steadily upward, depending on the configuration. Given the P6's size, power consumption, and heat dissipation (it requires active cooling), don't expect it to appear in battery-powered laptops soon.

The first customers—software developers and users of such applications as CAD, desktop publishing, scientific visualization, and statistics—are those who need as many MIPS as they can get. Server customers, however, march to a different beat. Those running NetWare file and print servers aren't CPU-bound and don't need the extra horsepower. But applications servers—NetWare-, OS/2-, Windows NT-, or Unix-based boxes running Lotus Notes, messaging engines, databases, or document repositories—are prime candidates for moving to the P6.

Early upgraders share another, somewhat surprising trait, says Eric Harslem, senior vice president of the product group at Dell Computer (Austin, TX): They tend to work

for smaller businesses, where systems are running off-the-shelf applications critical to the productivity of the whole organization. By comparison, larger MIS-directed organizations buy into new systems later, typically after one to three months of evaluation.

"[These companies] tend to have more custom apps and in-house-developed software, so they need a longer period to test compatibility," Harslem says. For example, if a vendor such as Dell decided to fit P6 systems with a new 128-bit graphics card to balance performance, it could break existing applications not written to standard APIs.

Your Basic P6

What will the typical P6 system be like? The processor will operate initially at 133 MHz, and the external CPU bus will run at one-quarter, one-third, or one-half of that clock rate, at the discretion of the systems designer. An Intel P6/PCI (Peripheral Component Interconnect) chip set called Orion will support version 2.1 of the 32-bit PCI bus at 33 MHz, but it won't support the 64-bit PCI extensions, says John Hyde, technical marketing manager for the P6 in Hillsboro, Oregon.

Due to the built-in level 2 cache, most P6 systems won't have external cache or cache controllers. Main memory will likely be conventional 60-nanosecond DRAM or, in some systems, EDO (Extended Data Out) DRAM, a new higher-speed alternative supported in the Intel Triton chip set for the Pentium. Memory configurations of 16 MB

can churn out as many as six micro-ops. That's better than the similar decoders in the Nx586 and K5 processors, which can generate a maximum of only four RISC86 instructions or R-ops per clock cycle, respectively.

Although Intel's micro-ops are comparable to an internal RISC-like instruction set, Intel says the P6 does not allow external access to the micro-ops. In other words, you can't write or compile a program in micro-ops and run it faster on the P6 by circumventing the decoder. The K5 has no provision for this, either. The Nx586 does permit such a bypass, although it's unlikely it will ever be used for anything but in-house testing by NexGen engineers. If Intel permitted this on the P6, it would raise the spectre of a new generation of software that discarded the troublesome x86 instruction set altogether.

After the decoder converts the x86 instructions into micro-ops, stages 7 and 8 finish preparing them for superscalar issue. In stage 7, references to x86 logical regis-

ters are reassigned to physical registers by the RAT (register alias table). This is similar to the register renaming found in other advanced microprocessors and is yet another way in which the P6 overcomes a historical limitation of the x86 architecture.

Ever since the 8086, the x86 architecture has made only eight integer-type GPRs (general-purpose registers) visible to programmers. That's a small register file in modern terms, but Intel can't expand the number of architectural GPRs without disrupting the x86 standard. Register renaming solves that problem by adding more physical registers to the chip and renaming them on the fly to represent logical (architectural) registers. In other words, at any given moment, any of the physical registers in the expanded register file can represent the logical registers that x86 programs expect to see.

The P6 has an impressive total of 40 extra GPRs, in addition to the architectural set of eight integer and eight floating-point GPRs. By comparison, the M1 has a total

of 32 GPRs, the K5 has 16, and the Nx586 has 14. The P6's expanded register file allows the CPU to explore more deeply when it speculates beyond predicted branches.

Furthermore, the 40 extra GPRs in the P6 are wide enough to handle floating-point values and integer values, as well as some additional status information that's appended in stage 8. The expanded register file functions as a general-purpose instruction pool that Intel calls an ROB (reorder buffer). This ROB is a 40-entry array of content-addressable memory that's arranged as a circular FIFO (first in, first out) buffer. Micro-ops (now a uniform 118 bits long) are held in this pool during various states of completion. Status bits record the state of each micro-op and provide additional information, such as which execution unit can handle that type of micro-op.

Deliberate Disorder

Until this point, the P6 has been processing the instructions in their original program order, just like the Pentium and earlier x86 chips. But this is where the P6 differs radically from the Pentium.

In-order CPUs frequently stall while they wait for a time-consuming instruction to complete. Because they must execute the instructions in the same order as they're arranged in the program, any instruction that causes a delay interferes with the smooth throughput of the pipeline—just as a factory assembly line halts when one worker takes too long to finish his task. To prevent this wasteful idling, the latest CPUs can execute instructions in almost any order they choose. Instead of waiting around for a poky instruction to complete, the CPU can fetch and execute a following instruction. Of course, to avert a software disaster, the CPU eventually stores the results in the order the programmer intended.

Some sequences of instructions must always execute in order because one instruction depends on the result of a previous instruction. This is called a *true dependency*, and the P6 is careful not to mix up these sequences. Among the status bits appended to each micro-op is a dependency flag and a time stamp that records the micro-op's original position in the instruction stream.

Out-of-order execution is a vital feature of the P6, Nx586, K5, M1, and many RISC chips. When combined with branch prediction and speculative execution, it allows the CPU to adapt the instruction stream to the availability of its internal resources. In effect, the CPU is its own optimized compiler. As long as there are

---

will be common, with a growing number of systems supplying 32 MB. Servers, and increasingly desktops, will ship with ECC (error checking and correcting) memory (10 bits per byte).

Initially, P6 systems will offer both PCI and EISA/ISA buses for performance and backward-compatibility. However, as the number of add-ins available for PCI increases, the need for EISA and eventually ISA will lessen. "Everybody in the world can hardly wait for EISA to go away," says Dell's Harslem.

A major factor here will be the growth of PCI-to-PCI bus bridges, which are supported in PCI 2.1. The biggest problem with the use of PCI today is its limit of 10 loads, which typically works out to only two to three add-in cards. Bus bridges, like network bridges, allow a larger number of devices within the same logical space while reducing traffic via filtering, says Digital's Lu.

Bridging multiple PCI buses will allow, for instance, systems with more useful configurations of six slots, obviating the need for other I/O buses. Having more available PCI loads will allow attachment of not just graphics and storage but also high-speed network interfaces (e.g., 100-Mbps Ethernet, FDDI, and ATM) and high-speed serial I/O

(e.g., Fibre Channel, IBM's Serial Storage Architecture, P1394 serial bus, and the Scalable Coherent Interface, or SCI).

Storage will likely be at least 730 MB, using IDE or SCSI. Most systems will ship with double-speed or faster CD-ROM drives. Basic graphics will probably have 1024- by 768-pixel resolution, driven by an accelerated graphics card with 2 to 4 MB of SDRAM (synchronous DRAM) or VRAM. A number of these systems will ship with NT, OS/2, Unix, or NetWare preinstalled. Servers and desktops will increasingly support SNMP or DMI.

More unusual configurations may include PCMCIA slots, quad-speed CD-ROMs, support for 40-MBps Ultra SCSI, built-in 10-/100-Mbps network ports, and on-board multimedia via DSPs (digital signal processors) or dedicated chips (wave table sound, video I/O, and compression/decompression). Some vendors may opt to use new memory types, 128-bit graphics accelerators, or the 64-bit bus extensions documented in the PCI specification. A few systems could sport new types of I/O ports, such as enhanced parallel or high-speed serial ports.

Andy Reinhardt is BYTE's West Coast bureau chief. You can reach him on the Internet or BIX at areinhardt@bix.com.

| Desktops Today ...	and Tomorrow
1994	**1996**
486/Pentium, 100 MIPS, 8 to 16 MB of RAM, VL/PCI bus, 540- to 730-MB hard drive, 10-Mbps Ethernet, sound, and graphics acceleration	P6, 250 MIPS or higher, 32 to 64 MB of RAM, 32-/64-bit PCI, 1-GB hard drive or higher, 100-Mbps Ethernet, multimedia via DSP, Plug and Play, voice navigation, digital video support, and videoconferencing

no dependencies, the CPU can continue predicting the outcome of branches and speculatively execute the instructions beyond those branches. There's no absolute limit to how far ahead the P6 can speculate, other than the capacities of its ROB, execution units, and other resources. Intel says the P6 typically looks 20 to 30 instructions beyond the program counter.

Branch prediction plays an important role in this process because the average program contains one branch every five instructions. (If you're a programmer and you're wondering why your own programs don't contain that many IF...THEN statements, remember that loop terminations and other statements compile into branches, too.) To reduce this number somewhat, the P6 adds one new instruction to the x86 instruction set: A conditional MOV that in some cases replaces a regular branch.

When a branch instruction enters the ROB, it is tagged with extra status bits that indicate a predicted target address and a fall-through address. Actual target addresses are stored in a 512-entry BTB (branch target buffer). The target address determines which cache lines will be retrieved during the next fetch cycle. Using a 4-bit dynamic history algorithm, the BTB keeps track of whether the branch was correctly predicted. This algorithm is smart enough to recognize patterns (e.g., alternating pairs of taken/not-taken branches) that would baffle some simpler

Larger than the average chip, the P6 is shown here at approximately 50 percent actual size.

algorithms, such as the binary flip-flop used by the K5. According to Intel, the P6's branch prediction is more than 90 percent accurate.

Thanks to its large BTB and ROB, the P6 typically predicts 10 to 15 nested branches, which is impressive indeed. Even some of the latest and most powerful RISC chips, such as the Mips R10000, can go only four branches deep.

But what happens if the P6 guesses wrong? This is when superpipelining works against the P6. A mispredicted branch can invalidate all the work in progress in previous pipeline stages—all the speculated instructions (and nested branches) beyond the mispredicted branch. The P6 must discard numerous entries in the BTB and the ROB. The emptied stages cause pipeline "bubbles" that impair the CPU's throughput until they are replenished.

There are two ways to measure the cost of a mispredicted branch: the actual loss of clock cycles required to recover from the wrong guess and the potential loss of work that was done speculatively and must be abandoned. For instance, say you're

The P7 and Beyond

Intel's P6 is the logical next step beyond the Pentium, but the P7 could introduce a radically different technology that achieves a breakthrough in performance while preserving backward compatibility. Then again, the P7 project could fall flat on its face.

Last year, Intel formed a much talked-about partnership with Hewlett-Packard to design a new microprocessor that is expected to appear in 1997 or 1998. The two companies are revealing little about this processor except that it will attempt to leapfrog RISC technology and run all existing software for Intel's x86 and HP's PA-RISC chips. In addition to supporting both of those legacy instruction sets, it will probably introduce a new instruction set of its own.

The prevailing rumor is that Intel and HP are experimenting with a technology called VLIW (very long instruction word). Ironically, VLIW is almost exactly the opposite of the technology used by the P6. While the P6 contains a sophisticated decoder that translates complex x86 instructions into shorter, simpler RISC-like operations, a VLIW processor would rely on a new type of compiler to pack several simple operations into very long instructions. Each instruction packet would contain operations that aren't interdependent, so the CPU

could rapidly execute them in parallel (see "VLIW Questions," November 1994).

In other words, a VLIW processor shifts the responsibility for instruction scheduling from the hardware to the software. The scheduling intelligence would be built into the compiler, which, in turn, would embed it into the applications software.

The compiler technology required to make this practical hasn't been perfected. Another problem is that software compiled for one version of a VLIW chip would probably have to be recompiled for the next generation. Software vendors would make millions on upgrades, but users wouldn't be too happy.

For these and other reasons, some observers doubt that Intel and HP can produce a commercially viable VLIW chip. Because the x86 market is far too important for Intel to bet everything on an unproven technology, it is likely that Intel has a parallel project to develop a more conventional P7 in case the VLIW project fails.

There is still plenty of room for improvement in the x86 architecture. The P7 could be a four- to six-way superscalar machine with larger primary caches, an integrated secondary cache, more execution units, and larger buffers to support deeper paths of speculative execution. It could also spec-

ulate both ways beyond a predicted branch, a technique IBM tried in a 1960s mainframe.

Meanwhile, Intel's competitors aren't sitting still. NexGen plans to introduce the Nx686 by the end of this year and predicts two to four times the performance of the Nx586. Cyrix says it is already working on successors to the M1, though no details have been disclosed.

AMD has mapped out its future in the most detail. The K5 will be followed by a K6 in 1996, with volume production in 1997. The K6 will be fabricated on a 0.35-micron process and is expected to have 6.5 million transistors; estimated performance is 300 SPECint92. In 1997, AMD plans to introduce a K7 that will ramp up to volume production in 1998 on a 0.18-micron process with 10 to 15 million transistors; it is expected to deliver 700 SPECint92 at 400 MHz. Finally, AMD is planning a K8 in 2001 that will have 20 million transistors and deliver 1000 SPECint92 at 600 MHz.

Additional competitors may surface, too. IBM Microelectronics, Texas Instruments, SGS-Thomson Microelectronics, and a few Asian companies already make 386 and 486 chips. So far, however, these companies have not committed themselves to the more difficult challenge of designing a state-of-the-art x86 processor that would compete directly against the latest products from Intel, AMD, Cyrix, and NexGen.

There are **Billions** of reasons to protect your software.

You only need **one** reason to **protect** with Sentinel:

It's the worldwide standard in software protection.

Piracy is the greatest threat to the world's software industry. Developers lose billions in sales to software piracy each year. Protect your software and get all the revenue you deserve.

More developers rely on Sentinel®, from Rainbow, than any other software protection in the world.

And for good reason. Sentinel performs where it matters most: leading the industry in technology, quality, reliability and support.

So when it's time to protect your DOS, OS/2, Windows, NT, Macintosh, UNIX, XENIX, LAN or CD-ROM application – protect with Sentinel. Protect with confidence.

Call the Rainbow office or authorized distributor near you for a Sentinel Developer's Kit and more information on why Sentinel is *Securing the future of software.*

ISO 9002 Certified

RAINBOW TECHNOLOGIES

SENTINEL *Securing the future of software*

playing roulette and bet a dollar on number 18. There's a 36-to-1 payoff if you win. But if you guess wrong, how much money have you lost? Your actual loss is the $1 you wagered, but your potential loss is the $36 you could have won by guessing right.

Similarly, the P6 pays two penalties for guessing wrong. Intel says the actual penalty is usually eight clock cycles, the time it takes to update the various buffers and registers, redirect the instruction pointer, and refill the pipeline. (A hinting mecha-

nism can reduce this to four clock cycles.) The potential penalty is more variable and depends on how many instructions were speculatively executed past the mispredicted branch. Because the P6 looks far ahead, this can easily be 20 or 30 instructions.

Intel says the P6, unlike the Pentium, is designed to deliver more of its potential speed without requiring optimized code. Nevertheless, there's no question that any compiler optimizations that min-

imize branches or enhance the CPU's ability to predict branches will boost performance—for all advanced CPUs, not just the P6. Intel says it is working with software vendors to develop postcompilation code profilers that will squeeze a little more speed out of the chips.

Superscalar Execution
Out-of-order execution in the P6 is managed by a unit called the *reservation station*, which is similar to the reservation

Smarter, More Powerful Servers

Among the most significant trends of the last two years have been the increasing use of x86-based systems as applications servers and Intel's growing role as a supplier of nonprocessor technologies, such as buses, networking, video compression, flash memory, and system management tools.

With the P6, Intel is continuing the trend of bringing high-end features to mass market. All the P6's internal registers are parity-checked, and the 64-bit path between the CPU core and level 2 cache uses ECC (error checking and correcting). Built-in diagnostic features, most new to P6, make it easier for vendors to design reliable systems: More than 100 events and variables inside the chip, such as cache misses, register contents, and occurrences of self-modifying code, can be counted and reported out via pins or software. Operating-system or utility software can read these values to gauge processor status and performance. The P6 also improves support for checkpointing (i.e., rolling back the machine to a known state in the event of an error), but again, the operating system has to be written to take advantage of machine-check interrupts.

The P6 also supports the same FRC (functional redundancy check) capability offered in the Pentium, in which two chips are lashed together to constantly verify each others' results and to signal an error if a conflict is found. Unfortunately, P6 doesn't solve FRC's main weakness—the nature of the error is undetermined.

Intel's P54C implementations of the Pentium introduced a simple and inexpensive approach to dual processing: a closely tied pairing in which host and slave processors shared a cache and divided up program threads transparently to applications.

(For more information, see "Pentium Chip's Dual Personality," December 1994 BYTE.) Only operating systems with support for multithreading were able to take advantage of this feature, so its market penetration has been low.

Better Multiprocessing
The P6 takes commodity multiprocessing to the next level, the Intel-defined MPS (Multi-Processor Specification) 1.1. Among the most difficult aspects of SMP (symmetric multiprocessing) is maintaining coherency among dedicated per-processor caches. Because the P6 handles level 2 cache coherency internally, its frontside (external) bus is inherently cache-coherent and presents, in effect, a kind of SMP bus to the outside world.

In the past, systems designers implementing SMP had to create their own buses to communicate among processors or license a solution such as Corollary's C-bus II. After claiming earlier that the 486 and the Pentium would be good for SMP systems, says Corollary's president George White, "this time Intel is right." The difference is the external bus. "You'll be able to tie four P6s together and make a pretty reasonable MP machine using Intel's cookbook," he says. But for now, that is the limit; the P6's arbitration logic supports only four CPUs.

Another dilemma facing vendors is that dedicated-cache multiprocessors typically benefit from using more than 256 KB of cache per processor. For now, Intel is limited to 256 KB in the P6 package. Thus, makers of high-performance servers who want to go beyond commodity SMP will have to use external cache controllers and SRAM (static RAM) to support more than

four CPUs or to implement larger caches. (The subtle trade-off, for which an answer isn't yet clear, is at what point a larger but external cache would surpass the performance of the P6's smaller, in-package full-speed cache.)

Intel can partly solve this problem with a bigger level 2 cache, which could be achieved by increasing the die size or moving to smaller process technology. But today its answer for vendors who want more than four CPUs is to closely couple, or cluster, systems across a high-speed memory-to-memory serial interconnect, such as SCI. Implementations of SCI for the PCI bus could ship as early as this year.

A Shot in the Arm
The combination of the P6 and MPS 1.1 will permit creation of a new class of "clone" servers that comply with a standard architecture and can run shrink-wrapped MPS-compatible operating systems. This could be a major shot in the arm for SMP. White cautions, however, that customers "may not be ready yet to buy servers on price alone." MPS doesn't go quite far enough, he says, in defining the kinds of features demanding users require, such as ECC, system management, and hot-swapping. These need to be implemented via hardware abstraction layers, or HALs (also known as processor-specific modules), the software interfaces added to operating systems to isolate them from hardware dependencies.

The irony for vendors is that if they have already written a HAL, they won't benefit much from MPS. On the other hand, customers will enjoy an attractive new range of options: From small vendors who lack the resources to develop HALs will come simple, low-cost SMPs, while high-end solutions customized to the machine will still be available at a premium price.   —A.R.

> You'll be able to tie four P6s together to make a pretty reasonable multiprocessing machine

stations in the K5 and the PowerPC. The reservation station is like a traffic cop; it's responsible for scheduling the order in which the micro-ops are dispatched from the ROB to the multiple execution units.

There are five execution units: two IUs (integer units), two AGUs (address-generation units) for load/store operations, and an FPU. The reservation station can dispatch up to five micro-ops to these units per clock cycle, but a sustained rate of three per clock cycle is more likely because some of the units are usually busy.

To schedule this dispatching, the reservation station checks the status bits of the micro-ops waiting in the ROB. If the micro-op isn't locked by a dependency, its operands are ready, and an execution unit that can handle that type of micro-op is available, then the reservation station dispatches it for execution. If more than one micro-op is ready for dispatching, the reservation station picks those that were generated in sequence. Although this pseudo-FIFO algorithm favors in-order execution, by now the instruction stream has become so rearranged by dependencies and branches that it's substantially out of order.

These steps occupy stages 9 and 10 in the pipeline, although in many cases stages 8 and 9 overlap. The next stage, execution, may consist of many stages if the micro-op is a relatively complex operation (e.g., a floating-point operation). Over the years, however, x86 processors have added much hard-wired logic, so the majority of integer operations (which make up most of the instructions in typical programs) execute in a single clock cycle, just like RISC instructions. Modern x86 compilers try to use as many of these fast instructions as possible.

To boost floating-point performance, the P6 has the same beefed-up FPU found in the Pentium, which means there is additional logic to handle multiply, divide, and shift operations. It's undoubtedly the fastest FPU on any x86 processor, even though it still falls well short of most RISC chips. Incidentally, the P6's FPU is free of the division bug that afflicts the Pentium. In fact, the FDIV bug was confirmed when Intel engineers testing the P6's FPU discovered that the P6's output matched that of the 486, but not the output of the Pentium.

Load and store operations in the P6 require one and two micro-ops, respectively. A load operation needs to know only the

memory location and the width of the data, so it's only one micro-op. A store must generate the memory address and the data; the decoder breaks the operation into two micro-ops so that they can execute in parallel. In fact, the AGU that handles stores has two ports and is separate from the load AGU, so the reservation station can dispatch a load and a store on the same clock cycle.

The P6 tries to maximize parallelism by spreading the execution units across five different ports on the reservation station. The two AGUs do not share the same ports, and neither do the two IUs. The FPU

Intel recently launched a new home page on the Internet's WWW (World Wide Web) that includes information about the P6. Among other things, you can read the technical paper that Intel presented at the IEEE International Solid-State Circuits Conference in San Francisco on February 16. Intel is also thinking about putting a P6 system on-line so that you can compare its performance to other machines. Intel's home page is at http://www.intel.com.

shares a port with one of the IUs, but because the FPU has such long latencies (due to the complexity of its operations), it rarely blocks the IU.

Time to Retire

After a micro-op has executed, its status flag is changed to indicate completion, and it is returned to the ROB. Thus, the ROB is a generalized pool for micro-ops that are awaiting either execution or retirement. Both the reservation station and the retire unit constantly scan this pool, looking for micro-ops with the appropriate status flags.

When the retire unit finds a completed micro-op, it verifies that the micro-op can indeed be retired. This is not as easy as it sounds, because the retire unit must put the micro-ops back into their original program order and watch for interrupts, traps, faults, breakpoints, and mispredicted branches. For example, an interrupt could invalidate any micro-ops that were completed speculatively after the one that was

interrupted. The P6 supports precise exception handling to the extent that the x86 architecture does, which means some instructions can be interrupted in midstream.

If a micro-op is ready for bed, the retire unit stores the results. This is when physical state is committed to logical state. If the original x86 instruction from which the micro-op is descended stores a result in a logical register, then the retire unit copies the result from the physical register that holds that value to the appropriate logical register. Intel refers to the logical register set as the RRF (retirement register file), but it's really just the architectural set of eight integer and eight floating-point GPRs found in any x86 chip. If the result of this retirement affects the status of another micro-op waiting in the ROB, the buffer is updated as well.

These steps occupy stages 12, 13, and 14 in the pipeline, although the last execution stage overlaps the first retirement stage. Retirement therefore takes two clock cycles, and the retire unit can commit the results of up to three micro-ops per clock cycle. In other words, the P6 is a three-way superscalar machine. The K5 is four-way superscalar, and the M1 is a two-way superscalar, like the Pentium, only more efficient. The Nx586 is harder to categorize because it can issue three or four RISC86 instructions simultaneously, depending on whether the optional FPU is present.

Ultimately, the final test of these competing microprocessors will be their performance when running everyday applications in typical systems. If the K5 and M1 live up to their promises, they should run neck and neck with the P6. The P6 and K5 appear to have more headroom for improved performance than do the M1 and the Nx586, but all four processors belong to roughly the same x86 generation—whether you call it five-and-a-half or six.

Intel commands such a dominant market share that it's unlikely the P6 will ever be seriously threatened by its rivals. And the P6 does a creditable job of defending Intel's technical leadership. But it's now clear that AMD, Cyrix, and NexGen no longer lag four or more years behind Intel with derivative clone designs. A new age has dawned, and users are the winners. ∎

Tom R. Halfhill is a BYTE senior news editor based in San Mateo, California. You can reach him on the Internet or BIX at thalfhill@ bix.com.

No other sound card ever made this kind of uproar.

"Creative Labs just keeps on improving on the design that made it famous. This latest revision of the Sound Blaster features a host of innovations. For example, it has both a 20-voice FM synthesizer and a 32-voice digital sample playback synthesizer for more realistic MIDI sounds. The QSound feature delivers surround-sound effects to your WAV files or you can choose to enhance them with such effects as stuttering and echo. Lots of useful applications come in the box, too, including the easy-to-use VoiceAssist speech-recognition software and well-rounded TextAssist for text-to-speech tasks."

HOME OFFICE COMPUTING / JANUARY 1995

"No ordinary sound card, the Sound Blaster AWE32 adds QSound to widen and enhance the stereo effects of sounds."

"The Sound Blaster AWE32 has Creative Labs' suite of sound playback and editing software that does the job with admirable simplicity--and the addition of sequencer Cakewalk Apprentice and the presentation program HSC Interactive SE rounds out the bundle nicely. And of course, the Sound Blaster AWE32 is compatible with software written for older Sound Blasters, so it will work on anything from Windows to Doom."

PC Computing/December '94

"I heartily recommend the ... nd Blaster AWE32. The 100% sound-effects compatibility and impressive music capabilities, make AWE32 an excellent all-round performer."

"To me, using a computer without sound and music is like driving a Lamborghini without the roar of its V12 engine. The

CD-ROM Multi Media Magazine November / December 1994

AWE32 is completely compatible with the whole range of Sound Blasters. The MIDI side of the AWE32 has been enhanced to provide it with the kind of sound you would expect from a professional synthesizer."

"...if sound quality and performance is important, you can't go wrong."

BEST BUYS

"For flexibility necessary for business-related audio, turn to the Creative Labs Sound Blaster AWE32, which stands out with decent sound quality and excellent software that includes 3D sound and text-to-speech and speech recognition capabilities. Its lavish documentation is another welcome feature. The AWE32 gave us the smoothest rides and the best sound in the games we tested."

174 PC WORLD • DECEMBER 1994

SOUND CARDS

"Creative Labs has long been the leader in this market, with its Sound Blaster series setting the standard for PC sound. The $399 AWE32 is the most advanced card in the Creative Labs lineup. It has everything you'd want in a general-purpose audio board."

Multimedia World/December '94

It seems the Sound Blaster™ AWE32 has the critics buzzing. For the first time, it puts awesome professional studio sound within the reach of the hungry audiophile. And of course it's genuine Sound Blaster compatible. So now that you've heard the experts' opinion, how about coming up with one of your own? Simply visit your local Creative Labs dealer. Or call 1-800-998-5227 Ext. 121.

CREATIVE
CREATIVE LABS, INC.

Circle 115 on Inquiry Card.

On-Road, On-Time, and On-Line

CHRISTINE WHITE

A load of fresh vegetables leaves California, destined for a packing plant in Little Rock, Arkansas. The driver, his load packed up and his route cleared with the dispatcher, heads out. He doesn't need to call the office until he arrives in Arkansas. While he's on the road, however, the shipper learns about a new, potentially lucrative job that's available if the trucker can change his destination from Little Rock to Hot Springs. Unfortunately, the driver won't know about the change until after he arrives at the original—and now less profitable—destination.

Such is life after deregulation for trucking companies, which count profitability by the mile. The more trips a driver makes

Jerry K. Watson, president and owner of Day & Night Transportation Services, with a Rockwell Pro2000 keypad.

TOM CASALINI © 1995

with a full load, the more money both the driver and the company stand to make. But a truck with a 50 percent load, or a driver sitting in a dispatch yard waiting for a run, doesn't make money. And a change in routing that wastes time and miles loses money.

Jerry K. Watson, president and owner of Day & Night Transportation Services, a trucking company based in Noblesville, Indiana, was looking to increase his firm's profitability and improve the service he could provide to his customers. The answer to both was to maximize each truck's drive time while minimizing the miles it traveled. As he investigated ways of achieving

these goals, he found a technological answer: a GPS (Global Positioning System).

Watson believes the Rockwell International Pro2000 GPS that his company now uses provides it with an edge over the competition. Day & Night's shippers are more likely to call the firm first and offer it its choice of loads. The company is able to attract and keep better drivers—who are paid by the mile for driving full, not empty, trucks. In addition, the company's dispatchers make better use of their time, assigning new loads rather than having to put out fires.

But GPS is an expensive technology to implement. Typical GPSes from Rockwell and Trimble Navigational, one of its competitors, cost between $2000 per vehicle for a receiver and antenna and $6000 for a system that includes communications software and other programs (e.g., for time-keeping and Department of Transportation logs). Dispatcher software ranges from $10,000 to $20,000, depending on the modules purchased. Added to that is the monthly fee for each truck's satellite transmissions.

But Watson is philosophical about these expenses. "How do you measure customer confidence? How can you put a price on the ability to better manage your people?," he asks.

Star Power

GPS is based on a constellation of 24 high-altitude satellites—officially known as Navstar GPS (short for Navigation System with Timing and Ranging). You can think of Navstar as the synthetic stars that make up a new celestial navigation system.

An ongoing U.S. government program, GPS has nearly $10 billion invested in it so far. Rockwell built the Navstar satellites, which were originally designed for the Department of Defense (DoD). System design began in the 1970s, and the last three satellites were put into orbit during the summer of 1993.

For more than a decade, these satellites were used exclusively by the government for various purposes. The DoD has now authorized limited civilian use of GPS. Three companies—Rockwell, Magellan Systems, and Trimble—have seized this opportunity by supplying GPS receivers and software to the trucking, automotive, nautical, aviation, and personal-navigation markets. Rockwell, which has been manufacturing receivers for government use since the

> **A trucking firm finds a combination of trucks, GPS technology, and computers pays off in fleet flexibility and control**

1970s, entered the consumer market in 1984.

GPS technology is a blooming industry, and new uses for the technology are now coming into play. Smaller vehicles, such as taxicabs, buses, and delivery vehicles, can make use of GPS applications. In addition, automobile manufacturers are starting to include GPS systems in their higher-end cars.

Trucks Meet Computers

Watson started the family-operated Day & Night trucking firm in the 1970s, when he bought his first tractor-trailer. Today the company has 60 trucks and employs 90 people. It makes regularly scheduled coast-to-coast runs to such destinations as Boston, Miami, and Riverside, California, although the majority of its runs cover 300 miles or less.

Five years ago, the company didn't even own a computer—which is not unusual in the trucking industry. Then in 1990, Watson, always looking for the competitive edge, bought an AS/400. Shortly afterward, he bought 386- and 486-based PCs running Microsoft Windows rather than AS/400 terminals.

During the past five years, Day & Night has been at the technology forefront in the trucking industry with its use of accounting software from McCormick Software and EDI (Electronic Data Interchange) programs from Sterling Software. The company added the GPS-based Rockwell Pro2000 Mobile Communications System in April 1994.

The deregulation of the trucking industry that took place in the 1980s created a crowded playing field. There are more trucks and trucking companies available than there are loads to carry. Add to that the skyrocketing costs of keeping a truck on the road—insurance, fuel, and fuel taxes, for instance—and it's becoming increasingly difficult to make a profit in this industry. Everyone's looking for an edge.

One edge comes from effectively managing drivers and equipment by maximizing drive time and minimizing miles. But with 60 trucks on the road at any given time, keeping track of their locations was a difficult job for Day & Night before it implemented its Pro2000 system. The dispatcher had to rely on the drivers' reporting their positions by telephone. That often meant there were large blocks of time during which dispatchers didn't know where a truck was. If a customer called to ask when a delivery would arrive, the only information the dispatcher could provide was the time the truck left the loading dock.

Day & Night's first solution to this problem was to use cellular phones. This enabled the dispatcher to call the drivers, and drivers could phone in when they needed to. This method was better than having no communication at all, but it proved to be expensive and unreliable. Costs for the phones were topping $100 a month per truck, and cellular service was not available in all areas. In addition, the service was sometimes abused by drivers making personal phone calls on the company dime. And there was still no way to verify a truck's location.

Watson knew he needed a system that would provide location information and communications without the drivers' being responsible for it. "I wanted the ability to look for a truck at any given time, to know where they were, to send them messages that I knew they'd read, and to better manage my resources," he says.

Then someone at McCormick Software told Watson about the Pro2000 system. Its dispatcher software runs on the AS/400 and integrates with the McCormick accounting software Day & Night

USING GPS FOR VEHICLE LOCATION AND COMMUNICATIONS

The Problems

- Half-empty trucks, dealing with new pickups, and accommodating changed plans.
- The home office's inability to find out exactly where trucks are.
- The need for the home office to communicate with drivers regarding new pickups and deliveries, new dates and times, changes in plans, and emergency information.
- The need for drivers to communicate with the home office regarding pickup and delivery status, problems, and delays.
- Providing security for drivers and their trucks.

The Solution

- The Rockwell International Pro2000 Mobile Communications System, which includes a GPS vehicle-location system and a satellite communications system.
- The Pro2000 system uses GPS satellite information and technology to locate vehicles anywhere on the earth's surface.
- Pro2000 also uses a geosynchronous satellite for communications with vehicles.

The Benefits

- A truck's location can be determined and relayed to the home office automatically, without any driver action required.
- The home office has accurate truck locations, allowing flexible pickups and deliveries.
- Wasted time is reduced, while truck availability is increased.
- Communications are less expensive than with cellular phones and are available everywhere.
- The system works with other software, such as accounting and mapping packages.

Lessons Learned

- Build on your existing computer hardware and software systems.
- Be aware of the costs of different levels of systems. Get a system that matches your current and expected needs.
- Look at a system in action before committing to using it.
- Expect initial resistance from drivers until they experience the benefits firsthand.
- Prepare to use the new flexibility to your advantage.

was already using. After several meetings with a Rockwell representative, Watson knew that the Pro2000 system would be able to meet his needs: nonintrusive positioning information, communications tracking, AS/400 compatibility, and integration with his company's existing accounting application.

Day & Night has outfitted about half its fleet with the Pro2000 system. All new trucks are delivered to the company with the system preinstalled, and Watson is now in the process of retrofitting the company's older vehicles that are still equipped with cellular phones.

The Pro2000 system consists of a transceiver unit, an antenna, and a keypad that resembles a PDA (personal digital assistant). The antenna, which looks like a small dome, mounts on top of the tractor's cab. It attaches to the transceiver, a small black box mounted under the truck's dashboard. The transceiver is used for receiving signals from the GPS satellite and for sending and receiving signals to the communications satellite. The keypad consists of a keyboard and a small LCD combined into one unit. The driver types outbound messages on the keyboard and receives incoming messages on the unit's screen.

Because the software for the mobile system is on a chip in the transceiver, only the antenna, the transceiver, and the cable that connects them need to be installed inside the truck. This installation, which is done by Day & Night's head mechanic, takes about a half day to complete.

A Pro2000 transceiver unit.

Rockwell spokesman Robert Woods reports that a Pro2000 system like Day & Night's costs about $4000 when purchased from Rockwell. Rockwell sells directly to owners of large fleets and uses dealers to handle small-fleet sales. These dealers might add a slight markup, and they typically charge extra for installing the hardware in trucks.

New Stars to Steer By
For complete and continuous global coverage, a GPS requires 21 satellites and three spares, all of which circle the earth once every 12 hours. The orbits, located about 12,000 miles above the earth, are

arranged in six overlapping orbital planes based on the earth's equatorial plane. These orbit patterns en-

A cutaway view of a Pro2000 antenna (inset) and an antenna mounted on a truck (below).

sure that any receiver near the earth's surface can receive signals from at least four satellites at any one time, although only three are needed for obtaining position information.

Each satellite has a receiver and a transmitter and constantly sends a unique, biphase, pseudo-random-noise code on two L-band carrier frequencies (1575.42 MHz and 1227.60 MHz). These codes are nothing more than random noise that's made up of a complex but repeated pattern of ones and zeros that has no intrinsic meaning or significance.

Each satellite's signal differs slightly from that of other satellites, and each sends two types of signals (hence the term *biphase*). One signal is *C/A code*, the signal received by all civilian receivers, such as those in the trucks. The other, *P code*, is reserved for military use.

A GPS transceiver located on earth, such as those in Day & Night's trucks, receives and decodes a signal, which includes the time that the signal left the satellite. Using triangulation techniques, the software in the transceiver can determine its precise longitude, latitude, and altitude (see the text box "How Triangulation Works" on page 64).

GPS technology can accurately calculate these measurements to within a range of 1 centimeter to 40 meters, depending on the type of transceiver used and the complexity of the software performing the triangulation. The system that Day & Night uses is accurate to within 25 feet, about

the width of an average suburban street. Once a truck's transceiver determines the position of the truck, Day & Night's Pro2000 communications software, which is called Exec2000, creates a positioning report that's sent to the dispatcher using a single geosynchronous communications satellite.

Individual trucks transmit their position reports to Day & Night's AS/400 at defined intervals. Determined by the dispatcher, an interval can be anywhere from minutes to hours for individual trucks and can be reset at any time and as often as necessary. The dispatcher can also determine a GPS position on demand.

A truck's location is represented on Day & Night's AS/400 by lines of text on a display that shows the number of miles it is from a waypoint; for example, "Truck 2322 is 9.5 miles from the Day & Night dispatching center in Noblesville, Indiana." Waypoints can be buildings, cities, exits on a highway, or any other physical location. The system manager can create new waypoints at any time.

Rockwell also offers Windows and OS/2 versions of the Exec2000 software that display maps in a graphical format and place pictures of the trucks on these maps at their actual locations. Both versions can also keep a log of a truck's progress over several hours, which is important when severe weather and road conditions exist. Because AS/400s and mainframes have terabytes of disk space, the amount of space required by a program to be installed on them usually isn't an issue. The PC versions of Exec2000 for Windows and OS/2 each require 16 MB of RAM and 50 MB of hard disk space.

Communications
While it's important for Day & Night to know a truck's precise location, the ability

The Pro2000 keypad has an LCD.

New Close-Up Offers Fastest Remote Windows Speed

Fastest Modem Communications

Now Close-Up, two-time winner of PC Magazine "Editors' Choice" award for remote communications software, has a brand new version that lets you view & control other PCs *by modem* — just as if you were there!

> *"Version 5.0 Close-Up was fast, but Version 6.0 has even better Windows performance and added features to keep it ahead of the pack."*
> – PC Magazine, July 1994

Why Is Close-Up So Fast?

Close-Up learns as it works. It uses AI (Artificial Intelligence) to compress all Windows video function calls. That's why Close-Up does so well in Windows Benchmark tests (see WinBench 95 chart).

Close-Up uses its revolutionary Photographic Memory™ so that once Close-Up has seen all or part of a Windows screen, it's memorized. Then as screens change, Close-Up only transmits new unmemorized data. Incredibly, with this technology *Close-Up gets faster & faster the longer you use it.*

Expert System Makes Communications Easy

You don't have to be a communications expert to get the results of one. Close-Up has an Expert System that automatically analyzes system components and configures your system for maximum speed.

WinBench™ 95 Ver 1.0 by Ziff-Davis Labs Tests Remote Windows™ Speed

The graph on the right shows the speed of the three leading remote control programs when transferring Windows screens. As you can see, Close-Up handles more pixels faster, meaning you spend less time waiting for Windows screens to transfer.

Non-Intrusive Technology

Unlike most remotes, Close-Up's installation does not modify any of your sensitive Windows files, including the SYSTEM.INI and WIN.INI files, allowing Close-Up and your PC to run at optimal levels.

Use Close-Up and Your Fax Software on One Modem

Save money by using just one modem with Close-Up *and* your fax software. Simply click the Close-Up sleep icon. What could be easier!

WinBench scores (average of 10 runs):
Close-Up	209,270
pcAnywhere	122,600
Reachout	101,900

WINDOWS & DOS

Close-Up saves time & money by allowing you to connect between PCs running DOS or Windows, eliminating the need to switch between products when connecting to different PCs.

pixels per second
- 200,000
- 150,000
- 100,000
- 50,000
- 0

Close-Up 6.0 · pc Anywhere for Windows 2.0 · Reachout 4.0

RISK FREE OFFER — Call 805-964-6767

Try Close-Up 6.0! If you're not absolutely convinced that Close-Up 6.0 is the fastest remote program, we'll give your money back! Order Close-Up *Host & Remote*, a complete system for only **$199**.

Upgrade your competitive communications program to new Close-Up 6.0 *Host & Remote* now for only **$69⁹⁵***.

COMPETITIVE $69⁹⁵ UPGRADE

NEW VERSION 6.0
Norton-Lambert
Close-Up
Modem Remote Communications

We invite you to reproduce this test. Modems: Hayes OPTIMA 288 V.FC. Computer: (Host) Gateway, (Remote) Micron, Pentiums 90Mhz, 32MB RAM, 128KB RAM cache, 500MB HD, IDE controller, no hardware disk cache. Video: (Host) STB Lightspeed 1MB VRAM, (Remote) Diamond Stealth 64 2MB VRAM. 800 x 600 pixels, 16 colors. SuperVGA.DRV 11/1/93. 75 Hz. MS-DOS 6.22, SMARTDRV 2MB cache. WinBench 95 is available on ZiffNet via CompuServe. All products are shipping versions. Test performed without verification by Ziff. WinBench trademark of Ziff-Davis Publishing Co. Windows trademark of Microsoft Corp. *60 day money-back guarantee if purchased directly from Norton-Lambert. Offer expires 9/30/95. All prices in US Dollars. Prices exclude shipping.

© 1995 Norton-Lambert Corp.

October 1994 Close-Up 6.0 · November 1993 Close-Up 5.0 · January 11, 1993 Close-Up 4.0 · February 25, 1992 Close-Up 4.0 · February 1992 Close-Up 4.0 · January 14, 1992 Close-Up 4.0 · November 18, 1991 Close-Up 4.0 · March 1991 Close-Up/LAN 2.0 · January 16, 1990 Close-Up/LAN 2.0 · October 1989 Close-Up 3.0A · October 31, 1989 Close-Up/LAN 2.0 · April 1989 Close-Up 3.0A · February 1989 Close-Up 3.0A · October 1988 Close-Up 3.0A

Circle 84 on Inquiry Card (RESELLERS: 85).

to communicate with it is also essential. In addition to displaying the results of the trucks' GPS location reports, the company's Pro2000 system serves as the communications command center between the dispatchers and the drivers.

Through the communications portion of the Pro2000 system and the Exec2000 software, dispatchers can send and receive messages to one vehicle or to an entire group. Exec2000 can also create message confirmations, such as when a message has been read by a driver, and track the status of messages as they change.

While GPS information and communications both use the same software and hardware, there's no connection between the GPS satellite system and the communications satellite system. Rockwell offers communications through the land-based Ardis network and also uses American Mobile Satellite Corp.'s geosynchronous satellite for transmitting messages. (Geosynchronous satellites, located 22,300 miles above the earth, complete one orbit per day and are thus apparently motion-

less with respect to the earth's surface.) AMSC has been designated by the FCC as the single-source supplier for L-band satellite transmission in the U.S. Its first satellite was to be launched in March.

The Pro2000 system requires a dedicated SNA (Systems Network Architecture) port to send and receive messages. Communications from the Pro2000 are sent through a phone line connected to this port. All messages are routed through Rockwell's base station in Washington, D.C, which is linked to the AMSC satellite.

How Triangulation Works

GPS (Global Positioning System) location technology uses a trigonometric formula known as *triangulation*. There are two methods of triangulation; one uses an unknown location, and the other uses a known location. Vehicle-location systems such as the one used by Day & Night use the unknown-location method.

A transceiver in a truck receives signals from three satellites. From these signals, the software in the transceiver

can determine the truck's position. The satellites orbit the earth at about 12,000 miles from its surface, but a satellite is rarely directly overhead. The varying positions of a satellite in its orbit create a variety of distances from the transceiver.

If a truck receives a signal from a satellite located 13,000 miles away, the location of the truck is narrowed down to a 13,000-mile sphere centered on that satellite. If another signal comes from

another satellite located 18,000 miles away, the location is narrowed further to the points where the 13,000-mile and 18,000-mile spheres intersect. Add a third satellite, located 16,000 miles away, and the area of intersection is narrowed to two points—one on the ground and one impossibly high above the earth.

The distance between a truck's transceiver and a satellite is calculated with the old "velocity times travel time" for-

Location, Location, Location

Communications satellite

Navstar satellites

Trucking company

Rockwell base station

Rockwell base station

A truck-mounted transceiver relays positioning information to three Navstar satellites that orbit the earth at an altitude of 12,000 miles. The satellites collate the positioning data to determine the truck's geographic location. A communications satellite links the truck and the company headquarters via the system provider's base station(s).

When Day & Night sends a message to a driver, the message travels through a long-distance phone carrier to Rockwell's base station. A computer located there identifies the sender and the intended receiver of the message via a data packet. Rockwell's computer then sends the message to the driver via the AMSC satellite and automatically bills the sender. Messages sent from drivers to Day & Night are routed in the same manner.

With the Pro2000, a dispatcher contacts a driver by sending a text message from the home system to the driver's keypad. The proprietary keypad's small keyboard has special function keys, such as Print and Save, and its LCD shows four lines of text at a time. Drivers access keypad functions by selecting from numbered menu items and read through long messages by pushing the scroll keys.

Drivers can also store messages and read them later. Directions to destinations, for example, are stored and recalled via an alphanumeric code. Rockwell will soon offer a transportable keypad that drivers can carry with them into truck stops.

An incoming message on a driver's keypad, which can consist of either predesigned or free-form text, sounds an alarm on the keypad. If a driver is on standby waiting for instructions, he or she can take a nap and know that the alarm will awaken him or her when instructions are sent.

Incoming messages can also be transmitted to a driver en route. In one case, a Day & Night driver left Miami on a Thursday with a load scheduled for delivery in Kentucky the following Monday. After leaving the shipper, the driver headed for a hotel for the weekend layover. When the destination suddenly changed the delivery date, the dispatcher sent the driver a text message instructing him to deliver on Saturday instead. The destination got the shipment when it wanted it, the driver got to go home for the weekend, and Day & Night saved the cost of a weekend hotel stay for the driver, as well as freeing up the truck for other possible deliveries.

While it's possible for a driver to send free-form messages from the keypad, Day & Night reduces its transmission costs by using forms. Resembling blank templates, the forms have predefined layouts and information areas. The driver, for example, might use a billing-information form and enter the pertinent data. He or she then sends only the data, which flows into a corresponding form on the AS/400.

Day & Night now uses satellites for all its communications. The satellites offer a substantial degree of speed and accuracy, although they cost more to use than land-based communications systems. After they are sent, most messages are received within 5 minutes. A confirmation lets a sender know when a message is received. The Pro2000 system also date- and time-stamps a message when the receiver reads it. This tracking feature is important to Watson, who points it out to drivers and dispatchers alike whenever he hears someone say, "I never saw the message."

Rockwell's charges for messages are based on the number of characters they contain. According to Woods, messaging typically costs a company anywhere from $45 to $55 per month for each truck that is equipped with a Pro2000 unit. Watson finds that Day & Night's messaging costs are slightly higher, but still only about half of what he pays for communications on his cellular-phone-equipped trucks.

Competitive Advantage

Day & Night has found the Pro2000 system invaluable in many instances. For day-to-day operations, Watson believes that just knowing where his equipment is gives him great peace of mind. With just the tap of a few keys, he knows where every outfitted truck in the fleet is located.

Moreover, there have been instances where the system has given Watson's firm a leg up on its competitors. For example, Day & Night regularly transports empty bottles from a supplier to a Heinz catsup bottling plant. The shipper uses three different transportation companies and divides the load equally among them.

Last November, a manager at the bottling plant called Day & Night to ask about the estimated delivery time for one of its trucks. The other two trucking firms had met with delays, and the bottling line was about to close down because it was almost out of bottles.

Day & Night's dispatcher checked the Pro2000 system and was able to visually determine the position of the firm's truck in relation to the bottling plant. The dispatcher reported that the driver was 20 miles away and would arrive in 30 minutes. The driver arrived within minutes of the time the dispatcher predicted. As a result, Day & Night was able to improve the confidence level of that customer.

Integration

Day & Night is beta-testing a new method of integrating its McCormick accounting software with its Rockwell Exec2000 software. Currently, a shipper sends an EDI document that provides instructions for picking up a load. The document contains such information as locations for pickup and delivery, contact information, type and amount of freight, and dates. This information creates an order in Day & Night's McCormick software system.

To integrate this process with the Exec-2000 software, the Pro2000 system pulls the information from the McCormick software and combines it with directions to both the pickup and drop-off locations. A dispatcher assigns a truck number, usually by looking at the GPS-calculated positions of the trucks, and transmits the information to the driver. "It's a better use of time," explains Watson. "The information

---

mula you may have learned in school. The satellite signal contains the time it was sent from the satellite.

To work correctly in a GPS, each clock in every transceiver has to be synchronized to the nanosecond. Each satellite has four atomic clocks: One is the working clock, one serves as a backup, and the other two are responsible for keeping the first two in sync.

When the transceiver receives a signal, the time included in the signal is immediately recorded. The GPS's software then calculates the difference between the send time and the current time to determine the lapse. For example, the transceiver might find that the lapse is 0.1 second. Using the aforementioned formula, velocity (186,000) × time traveled (0.1) equals the distance from the satellite (18,600 miles).

After determining the transceiver's distance from all three satellites, the software in a GPS can accurately determine the transceiver's latitude, longitude, and altitude. Then, based on the transceiver's unique identity number, it can transmit its position using a different geosynchronous satellite communications network.

How long does all this take? If a truck uses the same three satellites during the operation, it can take as little as 2 seconds. Otherwise, the calculation can take up to 30 seconds.

Triangulation is an old concept: Celestial navigation uses triangulation with the position of three stars. But it looks like triangulation will be around for years to come. Even the starships *Enterprise* and *Voyager* use triangulation to plot courses for strange, new worlds.

goes into the system only once. We get directions the first time we go somewhere and then have them for good. We don't have to ask the shipper to fax a map every time."

A driver usually receives a pickup order on the truck's keypad while on the road. He or she can then go directly from his or her original destination to pick up the new load. This helps Day & Night to provide faster service to its customers and keep more full trucks on the road.

With this integration method, a driver can send messages, such as the bill of lading information and delivery confirmation, back through the system. He or she enters the information required to complete the form and communicates it to the office. Then it flows through the Pro2000 system and into the accounting software.

Watson has a personal wish list for integration. He's looking forward to—and encouraging Rockwell to build—a document-imaging system. "Right now, we get confirmation of delivery only from the GPS system," he explains. "The driver still has the signed freight bill in his [or her] pocket. We have to wait for him [or her] to mail it, and we [must] receive it before we can bill the customer, because most of them want to see signed proof of delivery before paying."

Such a document-imaging system would include a small scanner in each truck. The AMSC satellite could transmit the document image directly to the Exec2000 software and then, via the integration with the accounting software, send the document as a binary file complete with the EDI invoice.

A View from the Road

Initially, Day & Night's drivers weren't exactly enthusiastic about using the GPS. They felt like Big Brother was watching and that they were being spied on by the office. But an unexpected event changed their minds.

A driver called in at 11:00 p.m. and wasn't due to call in again for another 24 hours. Within minutes after he hung up, his wife called the dispatcher with news of a family emergency. A few minutes later, the driver received a message instructing him to abandon the trailer and take the tractor home. Using GPS positioning reports, the dispatcher was able to locate another driver 60 miles away from the trailer and instruct him to pick it up. Thus, the GPS mobile communications system solved two problems: getting the driver home quickly and delivering the load.

An Exec2000 screen, displayed on Day & Night's AS/400, showing the current location, status, and availability of the company's trucks.

The Pro2000 system also offers drivers a degree of security. The system's communications capabilities allow a driver to notify the dispatcher in case of a breakdown, a flat tire, or other emergency. In case of an accident where a driver might be unable to send a message, a push of a distress button on the transceiver sends an emergency signal to Rockwell, which then forwards the signal to the Day & Night dispatcher. Once, when a Day & Night driver had an accident, the distress button was automatically activated when his truck rolled over. The driver used the keypad to notify the dispatcher that, while the truck was damaged, he was unhurt.

The Pro2000 system has also helped Day & Night to keep its trucks on the correct routes. Occasionally, dispatchers notice a truck that's slightly off route and notify the driver to make corrections so precious miles and time are not wasted.

GPS location reports are usually initiated from the truck's transceiver. However, the dispatcher can order a report at any time from the Pro2000 system. On one occasion, a driver was late in sending an expected report. Concerned about him, the dispatcher activated the GPS tracking signal and located the driver about 200 miles away from where he was supposed to be. This particular driver was attempting to take a vacation on company time and gas. He was unpleasantly surprised to find out that Day & Night would not be financing his trip.

All GPS systems, including the one used by Day & Night, have safeguards to ensure that they can't be tampered with or turned off. Since the transceiver is firmly mounted under the dashboard, it can't be easily disconnected. In addition, the GPS satellites constantly synchronize the clocks of the transceivers, so drivers can't alter time stamps on reports or messages.

When a tractor is stolen, the thieves are usually quite knowledgeable; one of the first things they do is forcibly remove the antenna and the transceiver. According to

Jeff Brady, Rockwell vice president of transportation systems, it is still possible to trace a disabled system by checking its last known position. The dispatcher can then alert local police, and the tractor is generally recovered.

Buying with Confidence

If you're looking to bring similar technology into your firm, Watson advises that you do your homework. You should first find the software you want to use and then buy hardware that can run it. Too many people do it the other way around. Systems such as Rockwell's Pro2000 make the most sense where there's a need for vehicle-location information as well as communications with remote vehicles. If your company already has computers, then it makes sense to find a system that's compatible with the existing systems.

When selecting a GPS, Watson recommends that you visit a site that already has one running so you can see it in a production environment first. He's pleased with his company's Pro2000 system, but, although he wasn't unpleasantly surprised by any of its features, he wishes he'd seen it in action first.

"It's like buying a new car," he explains. "The salesman doesn't say, 'Because this is a four-cylinder, when you step on the gas it doesn't just get up and go'; you've got to drive the car yourself to find that out. You might still buy the car, but you will know what to expect." After Watson got his company's system in place, some things weren't quite as fast and easy as he had expected, although he admits that his expectations may have been too high.

So, how did a small company like Day & Night, which didn't own *any* computers five years ago, become one of the leaders in GPS technology? Watson sums it up by saying, "Technology and GPS are the future, and you've got to be there. They're going to throw a party; whether you're there or not, the party will happen. I'm just getting to the party early." ∎

Christine White is a freelance writer based in Sacramento, California. She has written extensively about computers and technology in the transportation, manufacturing, and food-processing industries. You can reach her on the Internet or BIX at editors@bix .com or on CompuServe at 71075.2673@ compuserve.com.

How to Implement ISDN

**LENNY TROPIANO
AND DINAH MCNUTT**

Ah, ISDN: instant Internet connections; video-conferencing; efficient E-mail; remote file transfers. That's the promise. But the reality for many ISDN subscribers has its own litany: billing surprises; hardware incompatibilities; broken connections.

Fortunately, there's help. As ISDN becomes more widely available in the U.S., subscribers—as well as the phone companies—are becoming more savvy about the technology. This article presents some of the fundamentals you will need to know if you're interested in bringing ISDN into your organization.

First, installation and service costs vary widely, so cost-justifying ISDN takes some effort. We'll sketch out representative costs in the section titled "What Price Service?," but also see the text box "High-Speed Alternatives" on page 70 to decide if ISDN is for you. Second, determine your hardware needs and allow yourself some time to find out whether the equipment you have identified is compatible with that of your ISDN provider; hardware interoperability still isn't a given. Finally, arm yourself with an arsenal of questions for your phone company to make sure you get the services you require—but don't end up paying more than you have to once you get them (see "Dealing with the Phone Company" on page 68).

Digital Communications

First, we'll review the basics. Unlike communications with an analog modem, where digital information must be converted to an analog signal and sent over telephone lines, ISDN provides point-to-point digital communications. In addition, ISDN lines can handle both voice and data transmissions.

The most common ISDN service is BRI (Basic Rate Interface), which is composed of two B (bearer) channels and one D (delta) channel; this service is sometimes referred to as 2B+D.

Nuts-and-bolts advice for ordering and installing ISDN

Each B channel operates at 64 Kbps for all user traffic, including voice and data transmissions. The D channel, which operates at a speed of 16 Kbps, is used primarily for telephony signaling information. The D channel can also handle low-speed packet data, such as X.25.

Another type of ISDN service is PRI (Primary Rate Interface), which consists of 23 B channels and one D channel and has the same physical interface as a T1 circuit. (It's sometimes referred to as 23B+D.) In Europe, PRI has 30 B channels and one D channel (i.e., an E1 circuit). In both cases, the PRI D channel operates at 64 Kbps. PRI is primarily used for PBX solutions and larger-bandwidth requirements. This article will focus on BRI. (For additional information about ISDN, see "Other Avenues" on page 72).

What Price Service?

ISDN prices vary among service providers, but typical installation charges run as high as $600, and you may be required to pay the entire cost up front. Some RBOCs (Regional Bell Operating

Companies) offer reduced installation rates if you sign a long-term commitment. Just remember you might be billed for prorated installation charges if you cancel before fulfilling your commitment.

The other big charge you can expect to pay is for the actual service. In some areas, you might pay a flat-rate ISDN BRI cost of from $22 per month (plus off-hour connection time) to around $70 per month (for all-day service). If your local ISDN company meters rates, you also have to con-

cate one channel to data and use the other for voice or data by multiplexing the first B channel or by placing a separate data call to another location. Assuming you have a flat-rate tariff, the recurring monthly charge of a full BRI is just a little more than that for a standard business voice-grade line, and you effectively get two lines as a bonus.

If you're in an area where your RBOC meters your ISDN BRI service, you can take advantage of bandwidth-on-demand.

made to locations outside your local area. But if neither of these conditions applied, your charges would be just the monthly flat rate for all or part of that bandwidth.

What You'll Need

Installation and service costs aren't your only expenditures as you bring ISDN to your organization. Phone companies end their physical involvement with ISDN at the RJ-45 (i.e., 8-pin) modular connector that they install in your building. In the U.S., it's the responsibility of the customer to convert a two-wire (i.e., one-pair) U interface to a four-wire (i.e., two-pair) S/T interface. Most ISDN equipment connects to four-wire S/T interfaces.

Dealing with the Phone Company

Before you commit to ISDN, get answers to the following questions.

1. Does my location qualify for ISDN?

2. What type of switch is being used at the central office that services me?

3. Is the phone company up to National ISDN standards (e.g., NI-1 or NI-2)?

4. How many B-channel devices can I connect to my BRI service?

5. What is the entire installation charge? Is there a time-based commitment, and if so, what happens if I cancel prior to that?

6. Do you meter your usage? What charges apply?

7. Are there any discount plans, such as an after-hours flat rate, if you meter your ISDN service?

8. Which long-distance carriers can provide me with clear 64-Kbps-channel ISDN service?

9. Do you sell hardware for, or recommend hardware vendors for, my particular use?

10. Do you know of any ISDN users group or forum in my area?

11. What optional services can I subscribe to (e.g., caller-number ID or call waiting)?

12. How many directory numbers (i.e., telephone numbers) can I have? Can they be unlisted? If so, can this be done for free?

The standard converter used for this purpose is the NT-1 (network terminator). It's powered by electrical current and, in most cases, provides limited diagnostics functions to you and the RBOC. NT-1s range in cost from $150 for a simple model to about $800 for sophisticated units with such options as battery backup and analog POTS (plain old telephone system) capability.

Some ISDN hardware comes with an embedded NT-1, but this can be a limitation, because each U interface accommodates only one NT-1 converter. So, if you already have an NT-1 converter installed, you can't add a device with an embedded terminator, and you can add devices without terminators only if your NT-1 hardware has additional S/T digital outputs.

You'll learn that with ISDN there are no intermediate line conditions: Either it works or it doesn't. For this reason, we don't recommend daisy-chaining your internal cabling from one room to another. ISDN jacks may look similar to today's analog lines, but analog equipment plugged into an ISDN circuit will cause communications problems.

In addition, troubleshooting ISDN problems becomes more complex if you need to trace all the links in a daisy chain. It's much easier to remove one complete line that you know goes to only one device or jack and then methodically trace down the problem. Also, it's important to remember that ISDN jacks, if left open, tend to collect dust, which can short out the line.

For everything to work properly, the wiring inside your building should run straight to the NT-1 unit. You should use level 3 or higher twisted-pair wire for proper noise shielding. Although the National ISDN specification states that eight devices can

sider the charges for call setup (i.e., making the call) and connect-time charges. The latter usually ranges from 1 to 3 cents per minute, but those pennies can add up.

For instance, if you wanted to browse on the Internet using the WWW (World Wide Web), it would cost you $1.80 per hour (assuming you make just one call). If you spent 30 minutes a day doing this, it would cost you around $27 per month just for the ISDN BRI service, plus the charges you incurred with your Internet service provider. With intelligent hardware, you can bring your connection up and down as needed to curb connect-time charges.

One of ISDN's most attractive features is its ability to simultaneously transmit circuit-switched voice and data on a line by using both B channels on a BRI line. Using the appropriate hardware, you can dedi-

The call-setup and tear-down of the ISDN line take just a matter of milliseconds. For example, with two BRIs (64 Kbps × 4 = 256 Kbps), you can get the same effective bandwidth as you can with a fraction-256-Kbps T1 circuit (which your service provider does by multiplexing the T1 line). You bring up additional B channels only when the bandwidth requirements demand it, and you bring the channels down when they are not in use. Thus, you pay only for the connect charges for the amount of bandwidth that you need. By contrast, with a fraction-256-Kbps T1 circuit, you pay for the leased line whether you have anything going across it or not.

But it's important to remember that this example involves four separate calls; each would incur metered charges, if applicable, or long-distance charges if the calls were

communicate on a BRI, most RBOCs limit the number of devices to a maximum of two on each of the B channels. This avoids contention for these channels when devices are already using both B channels and another device tries to gain access.

Some vendors provide this functionality in hardware. For instance, when a data device is using both B channels and a voice call comes in (which is signaled on the D channel), some devices will drop one of the B channels (bringing the data connection down to 64 Kbps from 128 Kbps) and then ring the telephone for the voice call. The device reestablishes the data B channel when the voice call terminates.

National ISDN-1, or NI-1, and other ISDN standards give vendors a foundation for developing hardware that works with the various telephone switches at RBOC central offices. Without these standards—and sometimes even with them— interoperability is a problem. The common switches used today are the AT&T

High-Speed Alternatives

Analog modems are becoming faster—the top-end V.34-modem speed is 28.8 Kbps—and less expensive (street prices for V.34 modems are about $200, while basic ISDN hardware costs around $500). So, are analog modems a better choice than ISDN?

Your first consideration should be how you'll use your high-speed communications line. You may find V.34 modems the most cost-effective solution if you aren't doing a lot of interactive work over the line. For instance, V.34 should work fairly well for batch-oriented tasks, such as downloading E-mail or Usenet news articles. But it may not be the best solution for using a WWW (World Wide Web) browser.

If speed is your primary concern, consider that an ISDN connection can be up to six times faster than a V.34-modem link. For example, look at the theoretical raw transfer times for a 100-Kb file in the figure at right. (Note that many other factors, including protocol overhead, compression techniques, and data type, also affect transfer times.) Thirty-four seconds may not sound like a long time to transfer 100 Kb of data, but if you're using a WWW client, the difference between that and 12½ seconds seems significant.

Does fast connection time matter for your application? In addition to its fast data transfer rates, an ISDN line only takes a few milliseconds to bring up a connection. If you're running TCP/IP over your ISDN connection and your hardware supports on-demand connections, you can leave your connection down most of the time and bring it up only when you need it.

You might want to do this if you're paying for metered service for either your ISDN connection or your TCP/IP connection on the other end (or both). In any event, connection time for ISDN is negligible, which is not the case for analog modems, which must dial and negotiate the correct speed before the communications links connect.

High-Bandwidth Circuits
If you're communicating with just one service provider, you might consider a permanent 56-Kbps leased circuit, which compares closely in performance with ISDN. Compared to ISDN, leased circuits usually have higher installation costs and recurring costs. Recurring costs include charges for the mileage from your POP (point of presence) to the provider of your services.

By contrast, many RBOCs (Regional Bell Operating Companies) don't charge for mileage for ISDN switched circuits. In some areas, ISDN usage is metered, but many RBOCs offer either a flat-rate monthly charge (all day) or less expensive off-peak rates. Leased lines generally carry a flat rate, which can be several hundred dollars per month, depending on the speed of the circuit and the distance of your calls. However, if ISDN isn't available, or if it's metered and you are a high-usage customer, the leased-line flat rate might cost less. Also, many companies prefer to budget for a fixed recurring charge rather than gamble on a variable cost.

Hardware costs are similar for both ISDN and leased lines. With leased lines, you'll need a CSU (channel service unit) and a DSU (data service unit), as well as the appropriate hardware (e.g., routers, bridges, and concentrators) to transport the signal to the data equipment. The CSU and DSU handle A/D and D/A conversions, offer limited loopback capabilities, and provide the leased-line interface to the other DCE (data communications equipment). You can think of them as the "modem" for leased lines.

There is a break-even point where using leased lines becomes more cost-effective than using several ISDN BRI (Basic Rate Interface) lines because of the more sophisticated hardware required to do the inverse multiplexing of all the lines to gain the necessary bandwidth. For instance, assume you need a 512-Kbps uncompressed bandwidth. For ISDN, you'd need four BRIs (i.e., eight B channels). Start-up charges, which include installation for each of the four BRIs and an eight-port router that can do inverse multiplexing, come to about $6000. For the full bandwidth to be used, there have to be eight physically separate calls to the remote end, which also requires four BRIs and an eight-port router. If your RBOC meters your service, you'd be paying eight times the charges of a single B channel, since all the lines would be used.

To use leased lines, you would need a fractional T1 circuit with a 512-Kbps bandwidth. The start-up charges would include a single installation charge, a CSU/DSU (which costs approximately $1500), and a small router (which starts at $5000). The ongoing cost is a flat monthly rate for the fractional T1 circuit.

As you can see from these examples, in general, the higher the bandwidth required, the more attractive leased lines become.

100-Kb Transfer Times

V.34 modem (2.88 Kbps)
ISDN-1 (8 Kbps)
ISDN-2 (16 Kbps)

Seconds

No Other Client/Server Tool Manages Business Change and Complexity Like Magic

Integrating the Future

Managing the Present

Leveraging the Past

The development path you choose determines your company's success. With business practices changing and business environments being reengineered, you need client *and* server application development that stands the test of time. And that can only mean Magic. Magic leverages your legacy IT resources to provide an economical migration path to client/server. It manages your present development needs with an innovative programming paradigm that ensures extremely rapid application delivery and easy maintenance. And Magic integrates the future with fully scalable open architecture that takes applications from tactical to strategic - incorporating new platforms, DBMSs, multiple client/server models and application partitioning scenarios.

The secret of Magic's success is the Magic Engine, an advanced repository of sophisticated application functionality. Driven by a high level, table-driven interface, the Magic Engine lets the developer concentrate on business rules and logic instead of low-level repetitive tasks.

Isn't it time you put Magic to work for your business? For more information and to attend a free Magic seminar, call:

1-800-345-6244

MAGIC
Client and Server Development for Managing Change and Complexity

5ESS, the Northern Telecom DMS-100, and the Siemens-Stromberg-Carlson EWSD. It's a good idea to ask your telephone company which switch it uses and if it follows an ISDN standard. Then ask your hardware vendor if its ISDN equipment works with that switch.

Other Hardware

Depending on how you plan to use your ISDN line, you might need to install additional equipment beyond the basic hardware. Some issues to consider follow.

For voice-only applications, you can effectively consider the BRI to be two phone lines. Your existing phones won't work without intermediate hardware. At least one company, Adak Communications, sells TAs (terminal adapters) for around $800 to connect a POTS to one or both B channels with integrated NT-1s.

With a TA, your existing phones and wiring will all work properly. Just plug the phones, modems, or other analog telephony equipment into these units and dial normally. If your phone company doesn't normally do so, request that your B channels have two separate telephone numbers to distinguish the two lines.

ISDN telephone handsets are expensive ($500 and up), but they can provide limited PBX-type functions, such as call forwarding, hold, transfer, call waiting, conferencing, caller-number ID, and call blocking. But if you don't need this extra functionality, there's no point in paying the extra cost for these handsets.

If your ISDN needs are data only, your computer hardware will determine what ISDN equipment you'll use. Some vendors, including IBM, Silicon Graphics, and Sun Microsystems, provide BRI ports directly on their motherboards. If your computers don't include BRI ports, you have three connection options: an ISDN card that plugs into the computer's system bus, a high-speed RS-232 serial port, or a network-based adapter, such as a router or bridge.

The least expensive—and least flexible—of these options is a bus card. DigiBoard, IBM, and ISDN*tek are three vendors that sell such cards for prices starting at about $400. The most limiting factor of these cards is that they require special device drivers for the particular operating system

your computer uses. In addition, because they are internal to your computer, these cards usually provide little or no external diagnostics, such as LED indicators. Note that using such a card may require you to also buy an NT-1. Check with your vendor to see if NT-1 functionality is built into your bus card.

If you want simple high-speed connectivity for a single computer, a bus card is probably the best solution. Most cards come with software that has easy-to-use interfaces, and they don't require that you have a vast knowledge of hardware to get started. Since prices are always dropping, it rarely makes sense to pay for more functionality than you require. You might consider a bus-card solution today and plan on replacing it as the technology (or your requirements) surpasses your hardware.

A second connection method, which is usually a moderately priced solution, is the RS-232 serial port. Most computer systems have serial ports, but to be used for ISDN, a port must reliably handle rates of 57,600 baud for one B channel or 115,200 baud for two asynchronous B channels.

With an RS-232 port, ISDN equipment can act like standard analog modems in making a connection, and you can use standard communications software packages. No special drivers are necessary, and you can use existing serial-line protocols (e.g., SLIP and PPP) over the ISDN circuits.

One drawback is that you can get only 56 Kbps asynchronously out of the full 64-Kbps synchronous bandwidth on a single B channel. Also, CPU utilization is costly when serial ports are loaded down on fast asynchronous RS-232 connections. This may affect your computer's performance. Access-Works Communications, Adak, and Adtran offer serial ISDN modems in the $400-to-$800 price range.

The third, and most flexible, connection method is an ISDN bridge or router. Ascend, Combinet, and Gandalf Technologies are among the companies that offer such products. They connect to standard networks, such as Ethernet, and let you connect one or more local systems to remote systems via ISDN. You usually do not need additional drivers to use this method.

Price is one drawback, however: These devices start at around $1000. Also, installing the devices and configuring your systems require some networking knowledge. If your computer is not already on a network, you may also have to purchase a network card, cabling, and software.

A final note, no matter what method you choose to connect your ISDN network: Compatibility is still an issue, and it's wise to find out what type of equipment you'll be making connections with.

Compression Issues

Compression can offer 4-to-1 performance, or nearly 512-Kbps speeds for easily compressed data over two B channels. But not all compression methods are equal. No formal compression standard has been ratified, so many hardware vendors tout their own techniques. Unfortunately, this locks you into a particular hardware platform if you

---

Other Avenues

A number of resources are available for additional information about ISDN, including facts about hardware and standards. Here's a starting list.

● Service Availability
To find ISDN contacts at your RBOC (Regional Bell Operating Company) and information about ISDN availability in your area, call Bellcore's National ISDN hot line at (800) 992-4736 or connect to ftp info.bellcore .com in the /pub/ISDN directory.

● Tariffs
ISDN rates are generally set by the PUC (Public Utilities Commission) in your area, so if you have a complaint about rates, write to your PUC representative or attend the meetings at which rates are decided. Some areas support ISDN users groups, which can provide you with information or lobbying power when dealing with the PUC.

● Compatibility and Standards
One common method that hardware vendors use to improve interoperability is the PPP protocol. For more information, see "From Here to There," June 1994 BYTE, and the Internet Engineering Task Force's RFC (Request for Comments) 1618.

● ISDN Equipment
For more information about hardware, see the Usenet news group comp.dcom.isdn and the WWW (World Wide Web) pages at http://alumni.caltech.edu/dank/isdn/ and http://www.icus.com/isdn.html.

● For Further Reading
Frequently Asked Questions list from comp.dcom.isdn and ftp://rtfm.mit.edu/pub/usenet/news.answers/ isdn-faq.
Fritz, Jeffrey. "Clearing Away the ISDN Roadblocks." October 1994 BYTE. Also, "Digital Remote Access." September 1994 BYTE.
Kessler, Gary C. *ISDN: Concepts, Facilities, and Services.* 2d ed. Computer Communications Series, McGraw-Hill, 1993.

Color is color, unless it's brilliantly practical.

It's the business printer you've been waiting for. Professional color so reliable, inexpensive and easy to use, it's a practical office tool. Speed? It's the world's fastest color printer at 4 color pages per minute. Cost? Full text pages on plain paper at 3¢ each. Color for 11¢ on office papers. Simplicity? If you can load a stapler, you've mastered this machine. Reliability? Add

750 sheets and let it run overnight. That's robust. Cost? At $4,995, nothing in its performance range even comes close. Best of all, it's from Tektronix, a Fortune 500 leader, making world-class color printers for 13 years.

The Phaser 340 Color Printer. So practical, it's brilliant.

Tektronix

Circle 105 on Inquiry Card.

want to achieve the highest speeds.

Some vendors, including Ascend, are now licensing Stac Electronics' compression algorithm, which allows for interoperability among products and best-case compression of 4 to 1, according to Ascend. By contrast, Gandalf touts 8-to-1 compression using its proprietary algorithms. The choice depends on how important interoperability is to you.

Our best advice for choosing hardware is to shop around and try to gauge each hardware vendor's commitment to ISDN interoperability. Some service providers are also resellers for ISDN equipment, and they can usually pass along substantial discounts to their customers.

Phone-Company Talk

Talking about ISDN with your local telephone company can be a trying experience. Most of the larger RBOCs now have

ISDN BASIC-RATE-INTERFACE LINE TRANSLATIONS

If your hardware vendor can't supply the configuration information contained in these translations for the AT&T 5ESS telephone switch, consider another vendor.

PARAMETER	VALUE	NOTES
Terminal type	A	
Number of CSD	2	
Number of CSV	1	
Number of call appearances	1	
Display is Y/N	No	Note that the value of this field is not relevant for proper operation of the Ascend product.
Ringing/idle call appearances	Idle	This value is default for Terminal A.
Auto-hold is Y/N	No	This value is default for Terminal A.
One-touch is Y/N	No	This value is default for Terminal A.

ISDN and the Internet

The reasons for using ISDN to connect to the Internet are fundamental: cost savings, fast line speed, and high line quality. But you should first decide how you're going to connect to the Internet. You may need to locate an Internet provider that offers ISDN with the services you want at a price you can afford.

Like telephone companies, Internet providers charge either for a dedicated line (usually with unlimited use) or for connection time. You should estimate how much time you expect to be using your ISDN connection. Remember that the overhead of reconnections is low, so reconnecting each time you want to log on to the Internet is often an acceptable approach.

If you plan to connect a network of computers by using an ISDN bridge or router, find out if the provider can assign you a subnet of IP addresses or route traffic to a network of addresses that you already have assigned. There is no point of going to the expense of purchasing an ISDN router if you can't connect more than one host to the Internet provider.

Also, ask the ISDN provider for recommendations for hardware that's compatible with its equipment. This will ensure maximum performance and the fewest configuration problems.

vices might be connected to BRI supplementary services, such as call forwarding). The over-100 possible translations and the ability to mix and match scared away many first-time ISDN users.

Today, ISDN hardware vendors make it easier for customers to know which features to order by supporting a generic set of widely used options and by detailing the features that work with their products. For example, the table above, taken from *Technical Brief #5* from Ascend Communications, shows a list of BRI line translations for the AT&T switch. You provide the phone company with this information when ordering service. (Ask your ISDN hardware vendors for similar lists; if they can't provide one, consider different vendors.)

Before you invest in ISDN equipment, you first need to determine whether your building is qualified for ISDN service. Typically, you must be no more than 3.4 miles (i.e., 18,000 feet) from the central office. However, some RBOCs tout "ISDN anywhere" and are committed to provide service to even the most remote areas. The RBOC can use repeaters to boost the electrical signal or deliver ISDN to local neighborhoods via fiber optics (or multiple T1 dedicated circuits) by means of an IDLC (integrated digital-loop carrier).

Be aware, however, that if ISDN is new to your area, some of the technicians installing it may be doing it for the first time,

and problems can occur. A common installation problem is configuring the service with the wrong translations for the hardware you're using. Providing the technician with a list of settings similar to that shown in the table should prevent this problem. Sometimes engineers have trouble with the repeaters if you're outside the 3.4-mile range and aren't on SLC (subscriber-loop carrier) service. Installing repeaters usually means digging up roads and/or attaching them to telephone poles.

If your RBOC doesn't provide for ISDN service out of your local central office, some will, for a price, provide for backhauling it from another ISDN-capable central office. This FX (foreign exchange) method usually comes with additional recurring charges for mileage between the two central offices.

When dealing with your RBOC, it's important to determine what kind of service it provides. For calls out of your calling area, or LATA (local-area telephone authority), you also need to choose a long-distance carrier, called an IEC (interexchange carrier). Check to see whether they can provide you with end-to-end, digital, clear 64-Kbps connectivity to your destination. As always, long-distance charges apply even if you have a flat-rate plan from the RBOC. Some IECs are providing only 56-Kbps links once you leave the LATA, but that's still on the order of two to three times faster than what analog modems can provide. ■

Lenny Tropiano, a Unix software developer, is currently involved with ISDN and ATM networking for multimedia applications. He works for Applied Research Laboratories (Austin, TX) and can be reached on the Internet at lenny@icus.com/lenny/ or at the URL http://www.icus.com/lenny. Dinah McNutt is a freelance writer and consultant who specializes in Unix systems administration. She can be reached on BIX c/o "editors" or on the Internet at dinah@sysadmazon.com.

ISDN-speak in their vocabulary, but until recently, it was hard to find the right person to speak to. But once you find that person, ordering ISDN shouldn't be any more difficult than ordering a standard voice (i.e., analog) line.

Several years ago, to get the right kind of ISDN service, you had to determine your central-office-switch *translations* (e.g., how many directory numbers there were and how many B and D channel de-

Control Software Costs

**SALVATORE
SALAMONE**

When Household Credit Services (Salinas, CA) decided to install Microsoft Office throughout its network of 4300 nodes and 53 file servers, Microsoft suggested that it buy a per-seat licensing agreement. But before taking the plunge, Household Credit used Site-Meter, a software-metering program from McAfee Associates, to see how the existing copies of Office were used on its network, which it uses to manage several private-label credit cards.

"We found that concurrent usage of Microsoft Office rarely exceeded 10 percent," says Dale Gunderson, a communications analyst in charge of Household Credit's enterprise-wide metering project. "As a result of the findings, we decided to purchase a concurrent-usage license instead of the per-seat license." The result? By opting for this alternative licensing agreement, Household Credit expects to save $800,000 in licensing fees over the next two years, according to Gunderson.

Software-metering programs such as SiteMeter haven't traditionally performed such a dramatic role. Most companies meter software usage just to ensure they have enough legal copies of applications for all the employees who need them. Information like this is important defense against an audit by the Software Publishers Association (Washington, DC), or SPA.

But metering has recently become a tool used by companies to optimize software resources, which is becoming a necessity for companies with large networks like Household Credit's. Unwary corporations are finding that unnecessary fees, rising support costs, and redundant upgrades dwarf the initial software-licensing fees, which may already be too high.

Per-seat licensing is not the only problem. Because of poor inventory tracking, some companies buy additional licenses for

Corporations can reduce their licensing fees and support costs when they base their management strategies on actual usage

applications they already own. The Personal Computer Assets Management Institute (PCAMI, Rochester, NY), a group of vendors and users that studies the cost of managing hardware and software assets in organizations, estimates that public and private organizations in the U.S. last year spent as much as $2 billion for software they already own.

On another front, a Gartner Group (Stamford, CT) study found that software support accounts for 45 percent of the total cost of ownership of an application and that the handling of administrative tasks, such as ensuring that concurrent licensing agreements are enforced, represents another 13 percent. By contrast, licensing fees constitute only 14 percent.

Information Is Key

Savvy network administrators are successfully cutting the costs of software ownership by studying usage patterns to help guide

their software-purchasing and support strategies. What information is important? It can be as simple as knowing that only half the people in one division use a particular application. Or it could be more complicated, such as determining that one department is responsible for a disproportionate number of calls to a help desk with questions about using an application.

The software-metering and software-inventory programs that can help gather this information cost from about $295 for a single-server program to about $1200 for a 1000-user license (see the text box "A Class

of Their Own" below). Some vendors that currently offer these tools include Elan Computer Group (Mountain View, CA), Frye Computer Systems (Boston, MA), McAfee (Santa Clara, CA), Microsoft (Redmond, WA), Microsystems Software (Framingham, MA), Saber Software (Dallas, TX), Symantec (Santa Monica, CA), and Tally Systems (Hanover, NH).

Software-metering programs allow network managers to limit the number of users that can simultaneously access a particular application. If a company licenses Microsoft Word for 300 users, for exam-

ple, a metering program allows 300 people to run Word and restricts any other users until one of the 300 quits the application, thus freeing up a license. In addition, astute network administrators can exploit the more advanced features of metering programs to drive down software costs. Such features include logging usage by application and the ability to manage licenses so they can be shared (see the table "Metering-Tool Features").

Software-inventory programs offer a more passive form of license compliance by allowing managers to know what's run-

A Class of Their Own

As a class of products, metering programs offer remarkably similar features. Most allow a network administrator to automatically block access to applications once all the available licenses are used up, log the amount of time during which a user can control an application, set up classes of users so that some have a higher priority than others when accessing applications, and automatically queue up users who are waiting for a license to be freed up.

But not all metering packages are created equal. Depending on your network, one may be more appropriate for your needs than another. Items that you should consider when selecting a metering program include which NOS (network operating system) the program works with, which client applications it supports, and whether it's a stand-alone utility or one that can be used with other network utilities you might be running.

When it comes to NOS support, some packages, such as Saber Enterprise Applications Manager (or SEAM) from Saber Software, SiteMeter from McAfee, and Software Metering and Resource Tracking (or SMART) from Frye Computer Systems, run on Novell NetWare LANs. Others, such as CentaMeter from Tally Systems, Software Sentry from Microsystems Software, and Norton Administrator for Networks from Symantec, are NOS-independent: They all run on NetWare LANs, but they can also be used with other NOSes, including Banyan Vines, Pathworks, LAN Manager, LAN Server, LANtastic, and PC-NFS networks.

With client application support, all me-

tering programs support DOS and Windows applications. But some, such as Frye's SMART, McAfee's SiteMeter, and Saber's SEAM, can also meter OS/2 and Macintosh applications.

The Sum of the Parts
Companies that have turned to metering programs to help rein in the administrative costs of running a network often find they also need other utilities. For example, these companies might want to cut software-installation time by using an electronic software-distribution program. Or they may want to reduce the time it takes to perform an inventory by using automatic hardware- and software-inventory programs.

A company could buy stand-alone tools to do all these tasks. However, purchasing tools that are designed to work together often provides additional benefits. For example, integrated tools can let a network administrator display all network information on one management screen. Another benefit is that the data collected by tools designed to work together is often part of a common database and can be combined—in a single spreadsheet, for example—for easier analysis.

Several metering tools are actually one tool of many that fit into a plug-in-module approach to network-utility prod-

Graphical displays of application-usage levels, such as this one from Saber, make it easy for a network administrator to spot trends and adjust license purchases.

ucts. Most often, these tools work as stand-alone products, but they can also be used as part of an integrated management system. For example, Frye's SMART is one of several plug-in modules for the company's Frye Utilities for Networks. Other Frye utilities perform server management, automatic hardware inventory, network monitoring, and software distribution.

Similarly, SiteMeter from McAfee is part of BrightWorks, the company's integrated suite of network management tools, which includes programs for handling software distribution and hardware inventory. Microsoft, with its Systems Management Server, and Symantec, with its Norton Enterprise Framework, have taken this modular concept a step further and developed their own systems management architectures into which you can plug software-distribution, inventorying, and other network management tools.

ning on their network. Inventory programs produce lists of the applications that are residing on the network as well as on the local drives; these lists can be matched to software-purchasing invoices to demonstrate that the company is obeying the terms of its licensing agreements.

Paying Top Dollar

Software-metering and inventory tools can help a company maintain centralized control over software purchases. This lets an organization leverage volume-purchasing strategies that can lead to discounts of up to 60 percent, based on typical volume discounts and site-licensing agreements offered by applications vendors.

A software-inventory program can help a network administrator determine what applications are in use and in what quantities. "By knowing the real number of licenses that you need, you can negotiate price with vendors," says Bill Holder, director of Micropath (Bellevue, WA), an integrator that specializes in microcomputer-asset inventory services.

For example, Microsoft's Select program provides discounts on volume license and maintenance agreements for 1000 or more licenses. In April 1994, the company announced its Open License Pak, which offers discounts on as few as 50 licenses. Deals worked out under these programs include discounts on future upgrades, as well as provisions for phone support. Discounts, which are negotiated on a case-by-case basis, typically range from 10 percent to 50 percent, according to Microsoft.

Concurrent licensing can provide even greater savings in some instances. A corporation pays for a number of licenses that is less than the number of users on its network who will use the application. For example, a company might pay for 800 licenses even though 1000 users will use the program. However, only 800 of them can use the program at any one time.

The savings realized from moving from a one-to-one licensing basis to concurrent licensing is typically about 30 percent, according to the PCAMI. But others say that figure is low. "Conservatively, I think you can save 40 percent," says Holder.

Even if concurrent licensing saves 30 percent of the cost of buying individual licenses, that translates into a one-time $90,000 savings if a company with 1000 users buys a $300 program.

The savings can be even greater for programs that perform a special function, such as an organizational flowchart package or a mapping program that prints out directions between two locations. Such packages are often used by many people with-

in an organization, but only a few times a year. Seldom are there more than a handful of people running such a program at the same time. So, instead of buying 1000 copies of such a package, a company might be able to get away with a concurrent license for only six users.

One way for a company to plan for concurrent licensing is to first let its users run all the applications they need, without restricting access. Most metering programs let you run in this mode to gather data; they then visually display the information so it's easy to spot the patterns (see the screen in the text box). By examining this information, a network administrator can determine how many licenses are needed.

One note of caution: There's some debate in the industry about the liability risks involved when you collect information in this manner. Hard-liners in the licensing-agreement legal area say that allowing more users to access an application than you have valid licenses for is a violation of the licensing agreement, which makes your company subject to fines. Acknowledging this, some companies take a best guess as to the number of licenses they need: Once an application is up on the network, they use a metering program's data-collection ability to study usage trends and adjust the number of licenses for future purchases.

Nevertheless, the ability to examine trends in usage and then license accordingly saves money. The amount that can be saved varies by application, company, and even by department. For Household Credit, various departments have contributed in different ways to the company's projected $800,000 savings in licensing fees. "We estimated the initial savings of approximately $312 per user for customer-service personnel, who rarely use Office, and $31.25 per user for administrative personnel, who use Office more frequently," explains Gunderson.

Managing Users

With concurrent licensing, a company's cost savings can vary, depending on whether it licenses for the maximum number of users or the average number of users. Licensing based on average demand saves a company more money because fewer licenses are needed. However, when a company licenses for average demand, it will always encounter peak periods when some users are blocked from an application. Some companies find this situation unacceptable and always license for maximum demand.

But if a company chooses to license for average demand, many metering programs

can help. For instance, most display a message—with the relevant information, such as how many other users are waiting—on a blocked user's screen, telling the person he or she has been placed in a queue. When a license is freed up, the user is notified, and he or she has the option of either taking the license or passing.

In addition, many metering programs have an option that allows certain classes of users to always have access to licenses. For example, a network administrator can designate that 10 licenses out of 500 are to be used only by the company's 10 executives. Likewise, an administrator can ensure that the marketing department has more licenses for a presentation application than the engineering department does, even though users in both departments might frequently use the application.

License-hording is one sociological aspect of metering that network administrators must often deal with. Hording can be intentional, but often it's not: Many users simply load every application they'll ever need into their StartUp Group within Windows. This practice may seem logical to them; they may even think they're saving

METERING-TOOL FEATURES

Blocking
The key to all metering programs. Ensures compliance with terms of concurrent-licensing agreements.

Usage logging
Determines usage patterns on which to base purchasing and support decisions.

Cross-server license sharing
Dynamically reallocates underutilized licenses within an organization.

Queuing
Places users in a queue when all concurrent licenses are taken. Notifies a waiting user when a license becomes free.

Inactivity monitoring
Detects applications that are loaded on a user's desktop but go unused.

the company money by being more efficient—having the application a mouse-click away when they need it instead of having to wait for it to load each time they use it. But these users often don't realize that they might be blocking access to an application from another user who needs it.

To help solve such problems, some metering programs have inactivity trackers that can determine when a user simply loads a program and rarely uses it. For example, Tally's CentaMeter 1.0 uses agent software that intercepts common calls and detects every instance when a window is open. This helps CentaMeter to determine whether a program is being used or is simply sitting on a user's desktop as an icon.

Having knowledge about how long programs sit idle on users' desktops can help a network administrator negotiate with department heads when it comes to purchasing additional licenses. For example, one department might lobby for additional licenses of Excel because it can't tolerate having any of its personnel locked out of it. A preliminary check of the software-usage pattern might indicate that everyone in the department runs Excel when licenses are available, so the purchase of additional licenses may indeed be warranted. But with an inactivity tracker, a network administrator may find that some people are loading it and never using it.

Having such information enables sound decision-making. If software is merely sitting idle, loaded but unused, a network administrator might simply send out a memo explaining to users that this practice costs the company money. On the other hand, if all loaded applications are being used and many people are locked out of accessing a program, more licenses should be purchased.

Sharing for Profit

Metering software can also help you reallocate or share underutilized licenses. This results in the purchase of fewer licenses and can reduce your overall software expenditures by as much as 30 percent, according to the PCAMI, because many companies have software licenses that simply go unused.

For example, a company's accounting department might have 100 licenses for Lotus 1-2-3, but the most it ever uses is 75. Meanwhile, the company's expanding marketing department brings in 25 new employees and buys an additional 25 licenses of the same package. The company could have saved the cost of additional licenses by reassigning the 25 unused licenses from accounting to marketing.

Situations such as this can be avoided by using a software-inventory program that lets a network administrator know what software is sitting on the network and on local drives. With this information, an administrator could easily size up the situation and, in this particular example, save the cost of 25 licenses, or approximately $7500 (based on the $300 retail price of Lotus 1-2-3 release 5.0 for Windows).

Dynamic Reallocation

Besides saving money by identifying unused programs and shifting them to departments that need them, many metering programs allow for the dynamic reallocation of licenses between servers. This is different than simply finding unused licenses and making a one-time transfer to another department. With license sharing, licenses are transferred between groups on an as-needed basis.

For example, the aforementioned accounting department might use its 75 Lotus 1-2-3 licenses only at the end of the month to balance the company's books, but on a typical day use only a maximum of 60 licenses. By dynamically sharing the 15 licenses that merely sit idle for most of the month with another department, the company could save the cost of those 15 licenses.

Some companies have taken license-sharing to the extreme by dynamically reallocating licenses across different time zones (see "You're Saving Money when the Meter's Running," March BYTE).

Software Management Practices That Save Money

Reduce the number of licenses you purchase by not buying extra copies of software that you already own but sits unused on the network.

Centralize purchasing to reduce the cost per license by taking advantage of volume discounts and site licenses.

Opt for concurrent licensing agreements and use metering to ensure compliance with the terms of the agreements.

Dynamically reallocate or share licenses to get the most use out of each one.

Cut the time needed to troubleshoot software conflicts by using up-to-date inventory information.

Reduce recurring calls to a help desk by matching software-usage patterns to the frequency of calls to the help desk.

Charge back your support costs to individuals and departments based on usage, not on the number of people in a department.

Reduce network downtime due to virus infections by preventing unauthorized software from running on the network.

Avoid fines for using unlicensed software on company equipment.

Cutting-edge technology has no room for warm, fuzzy characters.

Soft, fuzzy type is great on kid's pajamas and schmaltzy greeting cards.

Unfortunately, it also shows up on bills of lading, insurance forms, and legal

briefs. That's why we've designed ACE Plus—Automatic Contrast Enhancement.

It turns soft-edge, low-contrast characters into cold, hard letterforms. And

that can actually make scanned documents more readable than the originals.

It also makes sure operators don't spend all day scanning documents over and

over. We have a brochure that tells you the whole story. For a free copy, call

1-708-675-7600. Or e-mail us: copiscan@mcs.com.

BELL+HOWELL

■ The Document Management People

Such cross-time-zone sharing is likely to be a short-lived phenomenon, however. One executive at a large operating-system development company, who declined to be identified, says that in the future, companies will probably be allowed to reallocate licenses between time zones, but they will likely have to pay twice as much per license to do so.

Most companies can save money by dynamically reallocating licenses within a smaller geographical area—and the cost savings are not limited to 1000-plus-node networks. For example, John Calmon, a programming-language administrator at the University of Pennsylvania in Philadelphia, administers a 25-workstation NetWare LAN and saves money by license sharing.

He originally purchased a metering program for license compliance, but he found that he could save money on license purchases by using information gathered by the metering programs in conjunction with the blocking feature once the licenses were purchased. He was able to cut the number of concurrent licenses for a word processing application from 16 to 10 (a 37.5 percent savings) and purchased three licenses (instead of six) for a $150 project management application—for a total savings of $450 for that application alone.

Treat Your Users Well

Cutting licensing fees certainly saves a company money. But these fees are not the biggest expense when it comes to using software; support and administrative costs can be much greater over the lifetime of an application on a network. In a survey of 180 large organizations, the Business Research Group (Newton, MA) found that companies spend, on average, $778 per user annually on Novell NetWare LANs. And, as noted earlier, a Gartner Group study found that software support accounts for 45 percent of the total cost of ownership of an application and that the handling of administrative tasks accounts for another 13 percent.

The use of metering and inventory programs can reduce a company's support and administrative costs. A company can cut its troubleshooting time by using the information gathered by such programs about what's running on a network. For instance, by quickly identifying (with an inventory program) that a caller to a help desk is using an old version of a program with an incorrect printer driver, a help-desk technician can save lots of time. Or a person might call to ask about a bug that has been fixed. A help-desk technician can quickly check to see if the caller has the

Safe Speed

RAIDM10   **RAIDM20**

5 Years warranty

SOLIDdisk RAID

Make sure your data RAID is fast and secure!

ou might not know how valuable your data is until it is no longer vailable, or even worse, it has been lost forever. If such an event were o occur, it will take valuable man-hours to restore the data, if possible, nd cost your business countless dollars.

his is why you should invest in a data secure, high availability system, efore it is too late!

OLIDdisk RAID High Performance Systems leaves no room for error. OLID COMPUTER offers a fully redundant fault-tolerant solution, as ell as the ability to have your data available continuously. This is ccomplished since all of the parts, hard disks, power supplies, ontrollers, and ventilators, are fully redundant, and can be exchanged

during operation in case a defect occurs. The **SOLIDdisk RAID** System is a true hot swap unit. There is no down-time, no tools required, and best of all, no cost to you!

There is only one way to guarantee data security, a **SOLIDdisk RAID** System!

Facts: Transfer rate 20 MB/s (100 MB/s in the future). Up to 80 GB/ unit. Max. 128 MB cache. Single or dual processor cache system. RAID levels 0, 1, 3 and 5. MTBDL more than 4 Mio. hours (500 years). Supported computer systems: **SUN SPARC, IBM RS/6000, Novell NetWare, HP 9000/800/900, Apple, SGI, Motorola, DEC-VAX, ALPHA.**

SOLID® COMPUTER
...the open systems solution pioneer!

Europe: Tel.+49-89-31 57 19-0 • Fax +49-89-3 15 91 46   **USA:** Tel.1-800•784•RAID • Fax 1-404 740 9761

Distributors: USA: DICKENS DATA SYSTEMS • Tel: 1-800-448-6177 • Fax: 1-404-442-7525 • **AUSTRIA:** SOLID COMPUTER • Tel: ++43-2231-5522-0 • Fax: ++43-2231-552226
SWITZERLAND: SOLID COMPUTER • Tel: ++41-56-701230 • Fax: ++41-56-713069 • **CZECH. REP.:** SOLID COMPUTER • Tel: ++42-2-436991 • Fax: ++42-2-434621
Circle 546 on Inquiry Card (RESELLERS: 547).

fix—a new driver, for instance—installed on his or her machine.

One network administrator at a Texas oil company says that he is able to cut, on average, one-third of the troubleshooting time from each phone call to his help desk by simply knowing what program (and what version of it) is running on a caller's desktop. Without an inventory program, many companies have trouble keeping track of who has the latest fix or patch. This makes it harder to provide support, prolongs the troubleshooting process, and adds unnecessary expenses.

The ability to track who has which patch or fix will become increasingly important as more users obtain access to on-line resources, such as America Online, CompuServe, and the Internet. Specifically, users will most likely see on-line notices that new patches are available and then simply download and install them themselves. Thus, a company's IS department will lose control of a process that it has traditionally managed. More important, an IS department without an inventory program won't be able to tell when a user has installed his or her own patch or fix.

Another way to cut support costs is to use information about software usage and become more proactive in managing a network. For example, a company might link software-usage patterns with the frequency of calls to a help desk from a particular department to plan training programs to reduce recurring calls. Or it could use information to develop a hardware-upgrade plan. For example, a company that wanted to replace the old Pentium chips in its systems could do so in a prudent manner: Instead of upgrading everyone in a day or week, the migration could be done gradually, leaving an IS staff sufficient time to handle its normal duties.

Moreover, an upgrade like this could be done in such a way that the people who needed the replacement the most got it first. For example, an engineering department that runs CAD/CAM programs would need a Pentium-chip replacement sooner than a marketing group using a desktop publishing program that doesn't rely on the Pentium's FPU.

Yet another way for a company to reduce its support costs is to use information about software usage to charge back the costs of support. With networks that have no centralized management, it's often hard to quantify support costs, and it is generally impossible to properly bill back the charges. But by using a metering program's logging functions, which can track how often a department uses a particular program, an IS department can charge back support costs based on usage. Metering programs allow equitable charge-backs for support. This charge-back concept is not a new one; it's similar to what was used in mainframe environments, where costs were charged back based on usage of processing resources.

It's not coincidental that software management strategies are in a sense reverting back to the methods used in mainframe days. For many years, networks have been allowed to proliferate without any centralized management. Now corporations are finding that the huge hidden costs of managing their resources, such as software on these networks, needs a centralized approach. Companies that have adopted such strategies are discovering the cost savings that stand to be gained. ∎

Salvatore Salamone is a BYTE news editor based in New York. You can reach him on the Internet or BIX at ssalamone@bix.com.

URGENT—YOUR INPUT NEEDED

On: WORKGROUP COMPUTING

Dear Reader:

To improve BYTE's coverage of technology in the State of the Art section, we'd like to get your feedback about what topics, areas, and products we should be considering, and in what ways. Later this year, we're planning on taking a look at workgroup computing. We intend to examine workflow software and groupware, programs like Lotus Notes and Collabra Share; the types of problems you can tackle best with current technologies; and the hardware and software we think you'll need in the future. But we want to hear your ideas, to find out about workgroup applications we may not be aware of, about interesting vertical-market products in this category. Also, we would like your help in identifying the people we should be talking to about this technology—users, vendors, researchers—you tell us. Finally, we're interested in hearing about what your organization is doing with workgroup computing—how you're using the technology and what problems and benefits you've encountered.

To let us know what you think, please use the following as a template to send us, via E-mail, an ASCII text file with your comments. Please be sure to include the <FIELDNAMES> with their angle brackets, followed by your information and comments. And thanks very much for your help.

Please E-mail the completed form to: **editors@bix.com**

```
<TOPIC>
Workgroup Computing

<LASTNAME>
Johnston

<FIRSTNAME>
Bobber

<TITLE>
Technical Specialist

<COMPANY>
Cybercyber Systems

<PHONE>
800-555-1234

<EMAIL>
bj@host.domain

<COMMENTS>
```
This is where your comments go. Be as brief or as long as you want. Tell us what you think, what you need, what you want to know more about. Tell us what you're doing. Tell us who we should be talking to.

Build Your Own WWW Server

BOB FRIESENHAHN

E ven if your company isn't hooked up to the Internet and has no interest in becoming part of the WWW (World Wide Web—also known as W3, or just "the Web"—a global network based on the Internet's TCP/IP protocols), there are significant reasons why you should consider using Web-based software and technology on your LAN. This article describes how to implement a Web server using readily available free software and equipment you might already own. It also points out pitfalls and obstacles to think about before beginning this type of project.

But you definitely should consider the benefits of weaving yourself into the Web. With HTTP (Hypertext Transport Protocol) servers and graphical client software, such as Mosaic and Netscape, you can achieve much of the information-sharing that's possible with Lotus Notes—but at a lower cost and without incurring the overhead of distributed databases.

Instead of replicating information to a local repository, the way Notes does, HTTP and Mosaic give you an efficient way of distributing up-to-date information and providing remote access to programs within your organization. With an attractive graphical interface, such as Mosaic, this results in a cost/benefit ratio that's hard to beat.

Many organizations want to set up their own World Wide Web server. Here's why—and how to do it.

Gopher, HTTP, NNTP (Network News Transport Protocol), Rlogin, Telnet, and WAIS (Wide Area Information Service).

WWW servers are designed to handle documents created using the HTML (Hypertext Markup Language) format. What makes HTML documents unique is their ability to include hypertext links that facilitate rapid access to other locations within the same document, to other documents on the same site, and to documents at another site entirely—the capability that makes the Web so powerful. HTML also supports internal references to external objects, such as image files.

The Web is based on the HTTP, which supports a client/server model. The protocol covers operations ranging from simple get commands to complex authentication mechanisms. HTTP is currently specified to run on top of TCP, the Internet's basic transport mechanism, but this is not an inherent limitation of the protocol.

To maximize the availability of the data on the Internet, many software packages—called Web browsers—have been developed.

How the Web Works

A major thrust of the WWW effort has been to unite existing protocols and formats into one common interface. To accomplish this, the URL (uniform resource locator) addressing format allows a user to specify any object on the Internet, along with sufficient information to retrieve it. Consider this example: ftp://ftp.uu.net:21/README. This URL says to retrieve the file named README from the site ftp.uu.net using FTP (file transfer protocol) on port 21. Commonly supported protocols and programs include FTP, Finger,

JAMES YANG © 1995

These packages use URLs to specify data location and can retrieve and display many forms of information. The most popular of these clients are Mosaic (originally developed at NCSA, the National Center for Supercomputing Applications, Champaign, IL) and Lynx.

Many variants of Mosaic are now available from such companies as Netscape Communications (formerly Mosaic Communications, Mountain View, CA), NCSA, Spyglass (Savoy, IL), and Spry (Seattle, WA) for popular platforms. All share a user-friendly GUI that can display text in different fonts and in-line graphics. Lynx is a character-based program for those who don't have the high-speed SLIP or PPP Internet connection required for effective use of GUI-based software. Usually, a Web client is thought of as a user-interface tool, but a WWW server can also be a Web client.

How Can I Serve You?

A WWW server delivers data via HTTP. A client opens a connection to the server, submits a single request, and receives the response; the connection is then closed. Only a small part of the URL—the path portion—is passed from the client to the server; the server already knows the rest.

When a full URL is passed from the client to the server, this indicates a *proxy* request that the server is expected to place on behalf of the client. Servers that support proxy capability act as clients on the requester's behalf. For example, an HTTP-only client can perform an FTP transfer by using HTTP to request an FTP URL from the proxy server. The proxy server uses FTP to retrieve the data and passes it back to the client via HTTP.

Proxy servers can also provide the benefit of a local disk cache, which speeds up the recovery of often-used URLs. Sites that use an Internet firewall can use a proxy server to allow ordinary clients to access the Internet through the firewall.

The WWW server is responsible for mapping a supplied URL into an object or responding with an error message. An object can take many forms. The simplest, and most common, object is an HTML text file. But it's possible for an object to be a program or to be built into a database request.

WWW servers are designed to mask implementation details from the user. Considerable flexibility in the configuration files allows the remapping of directory references into other directories, individual files, or programs.

Regardless of the type of object that the WWW server finally resolves to, the server is responsible for indicating the data format used for the reply. Response formats use MIME (Multipurpose Internet Mail Extensions) conventions. For instance, an HTML file would be indicated as type text/html in the response header.

Selecting the Server Platform

Deciding how—and whether—to implement a Web server, as well as what computer platform to use, depends on a number of factors. These include your finances, expected usage, data types to be stored, available hardware, and your experience

Server Directory Tree

```
/httpd ──┬── bin (HTTP server binaries)
         ├── cache (area used for proxy caching)
         └── server_root ──┬── bin (CGI server-side programs)
                           ├── config (configuration files)
                           ├── etc (access-control files)
                           ├── icons (icons used by server)
                           └── log (server log files)
```

This directory configuration is how I chose to structure the server files. Keeping everything together simplifies administration and future expansion.

with specific operating systems.

Finances. You must first decide whether you have sufficient resources and dedication to operate your own server. For a normal WWW server, you need a full-time Internet connection (56 Kbps or higher), a 486-class or better computer, software, physical space, and technically knowledgeable people. Or you can hire a commercial service to put your data on the Web.

Expected usage. If you expect a large number of accesses per unit of time—as a large organization or popular public-access site might require—then you should design your server accordingly. But if you are doing this only as a hobby, then a desktop PC with a good 28.8-Kbps modem and SLIP or PPP support might do.

Data types and quantity stored. Some forms of data, especially graphical images, occupy considerable storage space. Be sure your server has sufficient disk space.

CPU resources. Serving up text files on request takes up few CPU resources. However, you may choose not to limit your server to such simple tasks. For example, if your system or application stores data in a nonstandard format, it may need to run a program on the fly to generate a desired response in an accepted format. Similarly, if you're providing access to a database, your server will use significant CPU resources during database accesses.

Experience with operating systems.

Most people prefer to use operating systems that they're already familiar with. Luckily, WWW servers are available for many operating systems. The ideal operating system has excellent TCP/IP networking and efficient multitasking. This rules out MS-DOS and Windows. Unix, Windows NT, and Digital Equipment's VMS (with multithreading extensions) are all suitable. Unix is currently the best-supported operating system for WWW use and involves the least amount of risk.

Available hardware. Most sites connected directly to the Internet already use systems that a WWW server can run on. The main concern is whether your existing system (which may already support E-mail and Usenet news groups) has sufficient capacity to support a server. There are no clear-cut guidelines for choosing a server platform. A Pentium-based PC or Unix workstation (e.g., Sun's Sparc-Station 10) should provide good service in most cases.

There's more performance variance due to your operating system than to your hardware. A solid Unix implementation is an excellent choice. Failing that, VMS (for a Digital VAX minicomputer or an Alpha-based workstation) or Windows NT should be a workable solution.

Installing and Configuring Server Software

Once you decide to use a WWW server, you need software. (See the text box "Setting Up a Server," January BYTE, page 76, for on-line sources for WWW server software.) The example in this article is geared toward the CERN (European Laboratory for Particle Physics) httpd software, based on my experience.

Building a server from source code is often a good idea because it increases your familiarity with the software and ensures that the binary is 100 percent compatible with your system. But if you don't have an appropriate compiler or don't want to learn how to build a server, then a precompiled binary may be a better choice—and, with luck, everything will work fine.

There are few build-time choices because most options are configured at run time. With the CERN server, the only build-time option is to include SOCKS (a generic proxy implementation library) support. This allows the server to reside on a LAN that's inside a firewall and to access the Internet via a SOCKS proxy daemon that executes on the firewall machine.

If you're building a server to provide HTTP proxy service to clients on a LAN,

including this feature ensures the highest possible level of protection against unauthorized hackers on the network. Follow whatever building procedures are specified for the server software you select.

Before you install your server binaries, however, you should map out the server's directory structure. You must decide whether to install the binaries, configuration files, and log files where your system normally keeps such things or to put everything in a single directory tree. That's the approach I took, since having everything available in one place makes it easier to find related information and to move or copy the server to a new machine (see the figure "Server Directory Tree" on page 84).

Add an entry to your system's start-up script that automatically starts the server software when the system is booted—at least a command line that specifies the path to the primary executable as well as the configuration file to use. It's best to minimize the information specified on the command line; the configuration file is a more appropriate place for such data.

Most HTTP servers come with reasonable defaults and sample configuration files. It's best to start off small and add in features as your server grows. At a bare minimum, the configuration should specify the following information:

• The root (ServerRoot) of the directory tree (/httpd/server_root in the sample directory structure).
• The user identity under which the server will run. For security reasons (more on this in the next section), it's best to give the server as few privileges as possible.
• The path to the log files.
• A mapping to executables on the server. For example, in the CERN HTTP configuration file, the following entry will map any URL that starts with /cgi-bin/ (a common prefix) into the server-side bin directory and treat it as an executable.

Exec /cgi-bin/*/httpd/server_root/bin/*

• A mapping to HTML source files. If HTML source files are placed under a directory called /Web, the configuration entry to map URLs into that directory is Pass /* /Web/*. In the CERN server (and likely in others as well), the Exec and Pass mapping lines are matched in the order specified in the configuration file. So, to provide special treatment for a particular URL, position it before the default entries.

Security Issues

Current WWW servers provide simple authentication mechanisms similar to those used in FTP and Telnet applications (i.e., cleartext passwords matched against DES-encoded equivalents in a password file). This is poor security, but it's the current least common denominator between clients and servers (i.e., you can always use it).

More secure schemes based on Kerberos and RSA (Rivest-Shamir-Adleman) public-key encryption are currently under development. Netscape Communications has announced a secure client/server combination that uses RSA technology to perform authentication and encryption.

However, there's deep concern in the WWW community that marketing this technology without agreement by other WWW software suppliers may actually destabilize Web security standards. Recognizing this concern, Netscape Communications has joined the W30 Consortium to work out new security and HTML standards.

In addition to password-style authentication, access can be restricted based on network address—a domain address or an

This sample Mosaic forms application supports various queries of the Unix mail alias system.

**Aliaserv:
A Sample HTML Forms Application**

```
Aliaserv.htm
<html>
<head>
<title>Aliaserv</title>
</head>
<body>
<img src="aliaserv.gif" width="300"
  height="100" alt="">
<h2>Mail Alias Query Facility</h2>

<form action="/cgi-bin/aliaserv"
  method="POST">

Arguments <input size=30 maxlength=80
  name="arguments"><a
href="aliaserv-help.html"><img
  src="/Private/Images/Icons/
  action_help.gif" width="70"
  height="46"align="middle"
  alt="Help"></a><p>

<dl>
 <dd> <input type="radio" name="query_type" value="expa" checked>
  Expand Alias(es) Matching Argument(s)
 <dl>
  <dd> Select any or all of the following options for Expand Alias
  <dd> <input type="checkbox" name="query_options" value="-F"
   >Follow .forward files
  <dd> <input type="checkbox" name="query_options" value="-v"
   >Verbose alias expansion tracing
  <dd> <input type="checkbox" name="query_options" value="-V"
   >Verbose plus preserve name@localhost aliases
  <dd> <input type="checkbox" name="query_options" value="-l"
   >Long form. Repeat alias name and use commas as delimiters.
 </dl>
 <dd> <input type="radio" name="query_type" value="list"> List
   Aliases Matching Argument(s)
 <dd> <input type="radio" name="query_type" value="user"> List
   Aliases Containing Arguments(s)
</dl>

<INPUT TYPE="submit" VALUE="Submit Query"> <INPUT TYPE="reset"
  VALUE="Start Over">.
</form>
<a href="../main.html"><img src="/Private/Images/Icons/action_back.gif"
  width="69" height="40"
alt="MIS Page"></a> |
<a href="/"><img src="/Private/Images/Icons/action_home.gif"
width="69" height="40" alt="Home Page"></a>

<h5>Copyright Bob Friesenhahn (<a href="mailto:bfriesen@iphase.com">
  bfriesen@iphase.com</a>),1994</h5><br>
</body>
</html>
```

IP address mask. The CERN server supports groups, which can be individual user names, users at specified hosts, or other combinations. This capability, along with passwords, allows powerful access control.

Even if you don't plan to provide documents that require security, you should learn how to implement user-level security. If you offer access to a program and need to know who a user is, you must use authentication to acquire this information.

Currently, password-level authentication is the only available mechanism that can reliably identify a user on a TCP link. When authentication is required, the client displays a user-name/password challenge. The client caches the entry and passes it, in scrambled form, with each URL request sent to the same server (this ensures that the challenge is not repeated for each request). You can then use the authenticated user ID in executing server-side programs.

Design for Sharing and Growth

If anything remains constant, it's the need for change. It's far better to overdesign the directory structure of your server than to underdesign it. Use subdirectories and categorizations freely. Do this job well,

and few changes will be required as your server grows. This is especially important for hypertext documents; fixing broken cross-references can be time-consuming.

If your server will be developed and maintained by a number of people in your organization, take this into account. It's unlikely that you'll want to provide up-date privileges for all files by all users. So, split up your directory trees so they reflect your organization. This should result in maximum productivity while stepping on the fewest number of toes.

Some new conventions support WWW server development, including the concept of using a Webmaster and Docmasters. A *Webmaster* is responsible for the server's technical administration and overall structure. *Docmasters* are responsible for formulating and maintaining documents that reside on the server. In a shared environment, establish conventions that match your organization and your server goals.

Once you've built your server, you'll need to populate it with HTML documents and pages. As the listing on page 86 shows, formatting can be complicated. But a variety of editing and creation tools are available, including SoftQuad's HotMetal (avail-

able via anonymous ftp from ftp.ncsa.uiuc.edu:/Mosaic/contrib/SoftQuad), Cyberleaf from Interleaf (Waltham, MA), and AnchorPage from Iconovex (Bloomington, MN). Also, Microsoft recently announced an add-on to Word 6.0 that facilitates the creation of HTML documents.

Whither the Web?

The Web is a rapidly growing virtual document that comprises thousands of hypertext links to documents on sites around the globe. The HTTP server—the Web's underpinning—provides distributed access to documents, data, and programs. HTTP allows efficient sharing of information without the hassles and overhead of distributed databases. It also provides for platform-independent interfaces supported under Windows, the Mac, and Unix.

Even if your organization chooses not to participate in the WWW, you can use Web technologies on your LAN to provide common access to information. ∎

Bob Friesenhahn, a Dallas-based software engineer, moderates several BIX conferences. He can be reached on the Internet at bfriesen @simple.dallas.tx.us or on BIX as "thefuzz."

BARRICADING THE NET

As we venture into cyberspace, we need to make sure we're not leaving our own networks open to big-time trouble

Hey, hey! Let's be *careful* out there," cautioned Sergeant Phil Esterhaus after the daily squad room briefings on *Hill Street Blues*. We should remind ourselves of this when we turn on our computers.

As we computerize more functions, we must consider what's at risk, where the threats are coming from, and how to prevent them. We need to protect the confidentiality of information that's private, proprietary, and critical to our organizations' operations. To make good decisions on important matters—decisions whose consequences might be financial or matters of life and death—we must know that the data on which we base them is accurate.

Worry, Worry, Worry

We need to worry about numerous issues, including the following:

- Unauthorized persons accessing our data and our systems.
- Authorized persons causing damage, whether inadvertently or maliciously.
- Backups and storage.
- Virus attacks.
- Data transmissions being intercepted or eavesdropped.
- Security of financial resources and business transactions over the network.

With so much at risk, we must protect our data and networks with stronger mechanisms than mere passwords.

We addressed many of the security problems and solutions in "Distributed and Secure" (June 1994 BYTE). But as the volume of Internet use increases dramatically and as we connect our organizations' networks to thousands of other computer networks, serious security issues arise.

The firewall—a gateway through which all connections are made—offers a secure solution against intruders getting into and employees sending out information they shouldn't. Last October, *Internet firewall* was the buzzword of the day at the National Computer Security Conference cosponsored by the National Security Agency and the National Institute of Standards and Technology. You could approach a group of security professionals and say the magic phrase, and everyone would stop to hear more. Not everyone knew what a firewall was, but almost all believed they needed one.

Firewall Questions, with Answers

In this issue, we take a long, hard look at this hottest of security issues confronting today's information technologists. In "Build a Firewall," John Bryan examines the ins and outs of creating and installing firewalls. He looks at the architectural approaches for building a firewall, how they work, what problems they can cause, what factors to think about when planning a firewall implementation, and how much protection

you can expect from one.

In his companion piece, "Firewalls for Sale," Bryan looks at several types of commercial firewall products to show how they work, what the differences are between the various methods, and what you should consider when choosing protection for your networks. We also review one of the earliest-to-market firewall products, FireWall-1 from CheckPoint Software Technologies, on page 171.

The risks of working and doing business in cyberspace are present because there's also a tremendous potential for reward. As long as we keep our eyes open, assess the risks realistically, and take intelligent precautions, we can navigate cyberspace knowing that our own networks are safe from unwanted intrusion. ■

—Russell Kay, Technical Editor

Pollution Solution.

energy

EPA POLLUTION PREVENTER

This symbol identifies energy-efficient computers that save you money and reduce air pollution by "powering down" when not in use.

Over 2000 computers, monitors and printers have earned the right to bear the Environmental Protection Agency's Energy Star℠ logo. These computers are so smart, they go to sleep when not in use and wake up with a simple touch of the keyboard or mouse. Energy Star℠ Computers don't cost any extra, will save money on your energy bill, and prevent pollution from electricity generation. Be part of the solution and look for this logo when you make your next computer purchase. To receive a list of available products call the Energy Star℠ Hotline at 202 775-6650.

Buy It And Save.

BUILD A FIREWALL

Is your Internet connection putting your networks at risk? Establish a firewall and sleep easier.

JOHN BRYAN

The Internet is on fire. But as traffic increases dramatically, so, too, do the risks that your company's data may be sabotaged or stolen. Network firewalls have become a hot topic.

Relatively new creations in the computing world, Internet firewalls have their roots in control mechanisms and security measures that have long been standard practice in the mainframe community. But today's networked world has grown from the bottom up, with millions of new connections originating from personal computers and small networks. It's no longer possible to know who or what is on the other end of a network connection unless we take extraordinary measures.

What's in a Name?
A *firewall* is a barrier placed between your network and the outside world to prevent unwanted and potentially damaging intrusion of your network. Just as no physical fire wall is perfect protection against a fire, no digital firewall can make a network 100 percent secure against outside intrusion. But they can come remarkably close.

An important caveat to remember: Firewalls won't work, no matter how they are designed or implemented, without a clear security policy (see the text box "Network Security Starts with Workable Policy" on page 93). If a firewall is established for the wrong reasons, that can cause you problems, too.

Firewall Architectures
You can build firewalls in several ways, using a variety of mechanisms. The following examples are the most common:

• Router-based filters
• Host computer gateways, or *bastions*
• A separate, isolation network

The cost of a firewall can range from a $100,000 turnkey hardware/software

NIKOLAI PUNIN © 1995

system, installed and maintained by an outside vendor, to "no-cost" software available on the Internet from various suppliers and users groups. (Of course, creating your own firewall, even starting with free software, still requires a significant invest-

ment in time and people, which can quickly dwarf the cost of a ready-made solution.) One firewall-product vendor suggests that implementing a comprehensive firewall from scratch could require several worker-months, the equivalent of

perhaps $30,000 in salary and benefits. Rolling your own system carries with it other potential problems, particularly in maintenance and administration. You don't want a system that requires constant tweaking and expensive revisions.

Safe Network Services: FTP, DNS, and X11

One of the elements that binds the Internet together is a set of Unix-based programs—protocols, really—that let you access files and locate and use resources at remote sites. These services—FTP, DNS (Domain Name System), and X11 (the dominant windowing protocol for the Unix community)—are just now being augmented by the more automated operations of WWW (World Wide Web) browsers, such as NCSA (National Center for Supercomputing Applications) Mosaic and Netscape, but they pose serious security problems. Any organization considering implementing a firewall must be aware of how to make these services available yet secure.

FTP
After E-mail, FTP is probably the most-used service on the Internet today, letting millions of users download files and enabling organizations to make easily available a variety of documents, software, and graphics images.

But using FTP presents various security headaches. For starters, FTP isn't a single program; it's a protocol that's been implemented numerous times, and some of those implementations and FTP daemons are buggy. Another problem is that an FTP daemon normally runs with root privileges and doesn't relinquish this status. This means that if an intruder manages to get into your system by exploiting a hole in your FTP implementation, he or she can get straight to the heart of your system with maximum privileges. It was a known security hole in sendmail, another privileged program, that enabled Robert Morris's Internet worm of 1988 to propagate so widely.

So how do you get a secure FTP connection through a firewall? One way is to pass FTP requests to a proxy server—a restricted, cut-down version of FTP

that's known to be secure and that won't permit anything you don't want. A second approach is to modify the FTP software so that it will talk to only a restricted range of ports, for which firewall filtering or screening rules are in place.

Yet another approach is to use the pasv option, if it's available, to indicate that the remote FTP server should permit the client to initiate connections. This will work, however, only when the remote FTP server supports that operation. Finally, it is possible to build client versions of FTP that are linked against a SOCKS library (a generic proxy implementation).

DNS
DNS is vital to the correct operation of your internal mail and other functions—it's the telephone book that lets you know who's who and where in your network. But its structured database contains important and often sensitive information about your system that you don't want to make available to the outside world. Many organizations don't want to make their host names public.

For some of you, establishing an In-

ternet connection may require that you reconfigure the internal networks and domains. However, there are workarounds. One security system developer suggests an approach that works by redirecting DNS clients so that they talk to a DNS server that is on a different machine. First, on the outer side of your firewall host, you set up a DNS server that the outside world can talk to. To them, this is what your domains look like, although, in fact, they are seeing a restricted set of names (and aliases) that tells them only what you want them to know. We'll call this the public server.

Next, you set up another DNS server on an internal machine. This server is in fact the real thing, and it contains information about your hosts. You make sure that this server forwards any queries it can't resolve to the public server.

Finally, you set up your DNS clients—including any clients on the machine that hosts the public server—and use the internal server. This is the key.

In practice, an internal client makes a request and gets an answer from the internal server or—asking about an external host—the Internet but relayed

Public and Private DNS Servers

Private DNS server

Public DNS server

Internet

Internet workstation

- Internal DNS request
- Internal DNS request, name not known to private server
- External DNS request

Establishing dual DNS (Domain Name System) servers allows the internal network to keep its true domain names private. In the diagram, when an internal user makes a DNS request, it goes to the private DNS server, which returns the proper name. If the private server doesn't know that name, it passes it along to the public server, which then returns it to the user. All external DNS requests go to the public server alone, and it answers these requests directly.

Filtering Routers

Perhaps the simplest approach to creating a firewall involves using a programmable router—the type of device normally used to create a permanent, Internet connection to the outside world (often via a commercial Internet provider). Routers work by controlling traffic at the IP level, selectively passing or blocking data packets based on source/destination address or port information in the packet's header.

At the very least, you can use your router as a packet filter. This approach is probably the most common internetwork security mechanism used today. While reasonably good firewalls can be created with routers alone, it may prove difficult to program the router to exclude everything that you want kept out. Unfortunately, most routers come configured with a minimum of built-in protection, and many organizations simply install them this way.

The problem with the router-based approach stems from the variety of different protocols that are used on the Internet. At least three major network services are not handled well by packet filters—FTP, DNS (Domain Name System), and X11 all present special problems for the firewall implementor (see the text box "Safe Network Services: FTP, DNS, and X11").

through the public DNS server. External requests, however, are handled completely by the public server with its restricted information (see the figure "Public and Private DNS Servers").

While this technique can work well, it may not be the final answer. Hiding names in the DNS doesn't stop host names from "leaking" out in mail headers, signature files, or news articles. But this can be an effective way for an organization to quickly set up a restricted DNS gateway without having to conduct any significant internal reorganization or reconfiguration.

X11

X11, the predominant windowing protocol for Unix, assumes that the user's terminal is a server. While this philosophy enables many benefits, it also puts the system at risk. Applications connected to an X11 server have the power to seriously compromise security.

Spoofing—making your computer masquerade as a different machine or originating IP address, one that the receiving system trusts—is a problem that still hasn't gone away. In January, CERT (Computer Emergency Response Team) reported a new rash of IP spoofing incidents in which security software was stolen from the University of California's San Diego Supercomputer Center. According to CERT, some routers can be programmed to defeat such attacks, but others cannot.

Even though these are not X11-specific, they illustrate the dangers of spoofing. Remote systems that can gain or spoof access to a workstation's X11 display can monitor keystrokes that a user enters, download screen dumps that contain sensitive data, generate commands that appear to originate at the keyboard, and so on. Because of these problems, most firewalls block all X11 traffic unless they have a specially written application proxy to handle it.

Host-Based Firewalls

An alternative approach to firewall construction is to use a computer rather than a router. This offers many more capabilities, including the ability to log all the activity over the gateway. Indeed, when you think of a network firewall, you probably think first of a separate, highly secured computer system standing guard over your networks. This sentry system, sometimes called a bastion host, is a critical defense point that must be carefully designed, tightly controlled, and audited regularly.

While a router-based firewall monitors data packets at the IP level, hosts exert their control at an application level, where traffic can be examined more thoroughly. However, you can't use just any applications; you need to know that the applica-

Network Security Starts with Workable Policy

If you don't know where you're heading, it doesn't matter which direction you go. Similarly, no security measures will be effective unless you know what you want to protect. All computer security rests on the bedrock of a carefully drawn security policy that delineates what data and systems to protect, what levels of protection are appropriate, and what hardware and software you'll need to do it.

Policy sounds as if it's a high-level concept that doesn't have much to do with day-to-day operations and the nitty-gritty of firewall implementation details. But this is wrong and, what's worse, wrongheaded. Because any firewall system, no matter how it is built, can only implement specific decisions that have been made by the organization—as matters of policy.

BBN Internet Services (Cambridge, MA) identifies the "four P's" of Internet security:

- Paranoid—No Internet connection. Everything is forbidden, even perhaps what should be allowed.
- Prudent—Everything is forbidden except what is explicitly allowed.
- Permissive—The logical opposite of prudent. Everything is allowed except what is explicitly forbidden.
- Promiscuous—Everything is allowed, including those things that ought to be forbidden.

If your organization is like most, you'll find a comfortable spot somewhere in the prudent-to-permissive range and design your security accordingly. You need to determine what will and will not be permitted. Until your organization makes this fundamental decision, it has no workable basis for determining its security needs. The determination process, though, is far from simple and may well vary from department to department within a large organization. The fact is, security is one of those concerns that never really goes away. And in an age of connectivity and internetworking, security is more important than ever (see the text box "Top Security Threats" on page 94).

Finally, as conditions change, you must be ready to change or revise your policies. Creating your own "Ten Commandments" and engraving them in stone won't work in today's world.

tion software (and even the operating system) you run on this system may have its own gaping security holes.

To get around these problems and deal with potentially buggy protocols, host-based firewalls must use specialized software applications gateways and service proxies. These are, in essence, stripped-down versions of the original programs.

continued

Proxies and Host-Based Firewalls

In a typical host-based firewall, all Internet connections are handled through a single host system that runs proxy versions of such software as FTP, Telnet, Gopher, and other common programs. These proxies are written with deliberately limited capabilities so that access can be better controlled and the firewall made more secure.

Internal network

Host running only proxy versions of FTP, Telnet, and so on.

Filtering router (optional)

Internet

For instance, the standard versions of the Unix sendmail utility have perhaps 20,000 lines of code. A proxy version, such as Trusted Information Systems' (Glenwood, MD) smap (sendmail application proxy), contains only about 700 lines, because it doesn't include all the functionality of the standard version. It passes along mail messages only after verifying that they fit within the programmed restrictions. However, the cut-down nature of a proxy means that you can use it only with the application it's designed to serve (see the figure "Proxies and Host-Based Firewalls").

Isolation Networks

Another way to establish a firewall is similar to the host-based systems just described. Instead of interposing a host computer, you create another network, an isolated subnetwork that sits between the external and internal networks.

Typically, this network is configured so that both the Internet and the private network can access it, but traffic across the isolation network is blocked. Some isolation networks may contain only a single node configured as a bastion host that will support interactive sessions or application-level gateways. One advantage of isolation networks is that they can also simplify the establishment and enforcement of new Internet addresses, especially for large private networks that may otherwise face the prospect of having to undergo significant reconfiguration.

Roll Your Own or Hire a Mercenary?

If you choose not to buy a turnkey system, what other options do you have? You could certainly create your own if you have the expertise and required staff time. One valuable resource you should be aware of is the TIS Firewall Toolkit, created by

TOP SECURITY THREATS

A revealing look at what problems information security managers consider important—and what they're not doing about them

Half or more of those companies running mission-critical systems on LANs believe their security is unsatisfactory. That's one of the troubling conclusions contained in a 1994 year-end report by the auditing and management consulting firm of Ernst & Young. It surveyed 1271 information security managers to assess the current state of security practice.

The biggest problems cited are lack of resources—human (59 percent) and budget (55 percent). The biggest concerns are for network security and unauthorized external access (85 percent and 83 percent, respectively). In response to a different question, 93 percent expressed concern about the unavailability of network service, followed by fear of interference with operations, and loss of message confidentiality or integrity. A full 83 percent were also concerned about their inability to identify network users. Over half the organizations reported actual losses or interruptions in the past two years.

Connectivity shows up as a significant fact of life. For organizations that have over 2500 employees, 55 percent say their networks are accessed by customers, 46 percent by suppliers, and 33 percent by both. Some 45 percent use the Internet or other public data networks and 88 percent use E-mail. And even internally, most LANs and departmental minis are connected to a central computing resource.

According to the Ernst & Young study, the bigger the system, the safer it's considered to be. Only 4 percent of MVS mainframe users believed software security was inadequate. For Unix machines, the figure rose to 22 percent. LANs were in the 14 percent to 19 percent range, with NetWare at the top.

Desktop machines were considered the least secure, with MS-DOS and Macintosh computers reported as 57 percent and 47 percent, respectively. Interestingly, Windows (including Windows NT) did better, at 37 percent, and OS/2 systems were down at 27 percent, almost on a par with the Unix boxes.

The Concern and the Reality

OK, so that's what security managers say they worry about. But what are they doing about it? Here's a rundown of the extent to which they're using control measures:

- Antivirus software—91 percent
- Dial-back or secure modems—54 percent
- Firewalls—45 percent
- File encryption—36 percent
- PC hardware security devices— 33 percent
- Telecommunications encryption— 22 percent
- Message authentication coding— 17 percent

The actual use of security has not kept pace with the change in computing, especially the switch to interconnected networks. Let's face it, most of these controls have been around for a while; their use and importance are well understood, and the technology has been well developed. It doesn't take a rocket scientist to put them in place. We wonder what they're waiting for.

—Russell Kay, Technical Editor

Trusted Information Systems and available for no-cost downloading. (For more information, point your Web browser to this URL (uniform resource locator): http://www.tis.com/.) You can download bridge and host-based tools from Texas A&M University (ftp to net.tamu.edu, files in /pub/security/TAMU). Ohio State University offers a shareware version of Karl-Bridge, with limited features and hardware support (ftp to ftp.net.ohio-state.edu, files in /pub/kbridge).

Full-featured commercial versions of KarlBridge and KarlBrouter are also available from KarlNet (Columbus, OH, sales @karlnet.com). A version of Digital Equipment's "screennd" kernel screening software is available for BSD/386, Net-BSD, and BSDI.

If you don't have the in-house exper-

tise, you could hire a guru onto your staff or deal with an experienced consultant. But each of these will require significant effort and attention on your part.

One final approach to creating a firewall involves having someone else do everything for you—hardware, software, and administration. This may be especially attractive for the smaller company or even for the larger organization that simply doesn't want to commit people and time to the job. Outsourcing the firewall is simple and straightforward, and it means you don't have to worry about acquiring staff with the specialized knowledge needed.

FURTHER READING

For more detailed information on firewall and gateway systems, see *Firewalls and Internet Security: Repelling the Wily Hacker* by William R. Cheswick and Steven M. Bellovin (Addison-Wesley, 1994), which was reviewed in the September 1994 BYTE, page 42.

Firewalls is an Internet mailing list for firewall administrators and implementors. To subscribe, send "subscribe firewalls" in the body of a message to majordomo@ greatcircle.com.

Another helpful paper is "Thinking About Firewalls," by Marcus Ranum, Proceedings of the Second International Conference on Systems and Network Security and Management; available via ftp from ftp.tis.com: /pub/firewalls/firewall.ps.Z.

Dreams and Drawbacks
The Internet beckons us in some alluring ways. It promises a great deal in the way of rewards and benefits—connections with a multitude of individuals and organizations and access to information and resources on a scale heretofore unparalleled. And yet hooking up to the Internet can also be the source of significant dangers and risks.

Security is sometimes an elusive goal and can seem unattainable, especially when you think in terms of the exposure that an Internet connection offers. But there are workable, practical solutions on the market today. ∎

John Bryan is a freelance technology writer and consultant who is based in San Jose, California. You can contact him on the Internet at 5051339@mcimail.com or on BIX c/o "editors."

Circle 71 on Inquiry Card (RESELLERS: 72).

Never mind the box.
The goodies are inside.

When you buy a box of sweets, you couldn't care less about the box. It's the goodies you're after.

In the same way, the hardware in a UPS system is a necessity to make the system work in the first place.

But what turns this power backup into a true network power protection solution is intelligent software. Such as Fiskars' LanSafe III for Novell NetWare, Novell NMS, OS/2, UNIX and Windows NT environments. LanSafe III for Novell NMS is fully SNMP compatible.

Network shutdown

LanSafe III is the core of Fiskars' new network power protection solution. In the case of a power outage, it provides sequenced shutdown for the whole network.
Step 1: LanSafe III tells the workstations to send all data from their RAM caches to the server. Then it shuts the workstations down. It even saves all unsaved data in a Windows workstation and shuts it down in an orderly fashion - something no other software can do!
Step 2: Bridges and other communications devices stay on long enough for the data to pass through. Then they are shut down.
Step 3: The servers are supplied with power long enough so that all data is written on hard disk before the server is powered down.

This new level of security is made possible by LanSafe III and a wide range of PowerRite Plus and PowerServer UPSs, covering all your power needs from 250 VA to 6 kVA.

When other UPS vendors offer "easy battery replacement", Fiskars UPSs feature Advanced Battery Management (ABM™) technology which more than doubles battery life, typically outlasting the protected equipment!* And unlike competitive brands, Fiskars UPSs never crash the network even if a battery test fails.

Fiskars Corporation's UPS Group comprises two divisions: Deltec in the United States and Fiskars Power Systems in Europe. The Group has production facilities in the US, Finland and Mexico.

The quickest way to really protect your network is to contact a Fiskars authorised reseller right away. Phone us for the address of your nearest stockist.

Real-time graphical management console for complete UPS control on the LAN. Any UPS in the network is just a click away.

* In a traditional UPS the battery typically needs replacing every 2 to 4 years, while in a Fiskars UPS with ABM the batteries live up to 8 —10 years.

FISKARS

▌▌ POWER SYSTEMS

Fiskars Power Systems, P.O. Box 54, FIN-02921 Espoo, Finland
Phone +358 0 452661, fax +358 0 45266396

Circle 576 on Inquiry Card.

FIREWALLS FOR SALE

A look at five different firewall products and services you can install today

JOHN BRYAN

Sometimes, simply foiling an outside attack isn't enough. Nobody enjoys unwanted visitors constantly knocking on the door. One high-powered deterrent is Sidewinder, a $30,000 firewall system from Secure Computing Corp. (Roseville, MN). The company advertises Sidewinder as "security that strikes back."

SCC's background has been in developing security for U.S. government classified systems and networks, and the company is parlaying this experience in the commercial arena. Its client list still includes several "initialed" agencies. It offers hardware, software, cryptography, and consulting services.

Sidewinder is basically an outgrowth of SCC's Secure Network Server, a class A-1 (this is the highest rating awarded by the National Computer Security Center) server designed primarily for military installations. Sidewinder is hardware and software; it's designed around a 90-MHz Pentium PC and an SCC-tweaked version of BSD Unix.

SCC has modified this operating system so that it is secure in and of itself, requiring no proxies or gateway applications. In addition, Sidewinder is furnished only as a complete turnkey system, not as software alone. SCC has carefully examined the BIOS of the particular Pentium PC it uses and has modified its own software to make sure that there are no surprises or security holes to exploit.

Type Enforcement is what SCC calls the patented mechanism wherein the operating system and its applications stay secure. Under Type Enforcement, data and processes are assigned to class types according to your security policy. A central, protected table is used to enforce how any class may interact with another; this, too, is determined by your security policy (see the figure "Sidewinder's Type Enforcement").

Type Enforcement provides what SCC

NIKOLAI PUNIN © 1995

Sidewinder's Type Enforcement

Outside network

Internal network

Outside process

Outside program text

RE

RW

Unfiltered data

R

Filter process

RE

Filter program text

RW

Local process

RE

Local program text

Filtered data

R: Read only
RW: Read and write
RE: Read and execute
■ No access allowed

Domain definition table

	Local	Filter	Outside
Unfil		R	RW
Fil	R	RW	
Lcl PT	RE		
Fltr PT		RE	
Out PT			RE

Data and processes are assigned to classes, while a protected central data structure (called the DDT, or Domain Definition Table) enforces the ways each given class of process can touch a given class of data or programs. Whenever a data access is initiated, the DDT is consulted and a defined access value is assigned. In this figure, an outside process can execute its own program and can read and write unfiltered data, but it can't execute any other program or touch any other class of data. Data from the internal network can only reach the Internet after passing through a filter process.

calls *defense in depth*—meaning that, even if a determined hacker were able to break into the Sidewinder platform itself, he or she would be left stranded in one domain without access to any other applications or processes. And breaking in is made more difficult because Sidewinder can filter any data that passes the network boundary, coming in or going out, on the basis of its content and its source or destination headers.

Secure as Sidewinder is, that's not even its most interesting feature. Unlike most other systems, Sidewinder can provide active security—the part that "strikes back." Sidewinder can be configured to respond to any probe or attack, however slight, in a variety of ways.

When Sidewinder detects a hacker, the first thing it does is send a silent alarm to the system administrator. What happens after that is up to the administrator and company policy. The system can let the intruder in and permit certain activities up to a point, all the while collecting information on the source of the probe and what types of actions the hacker takes. The system can also provide dummy password files, dead-end traps, and other stealthy

defenses—a veritable "hall of mirrors," where nothing is quite the way it appears.

The intruder might, for example, issue a command to delete all files; a subsequent directory check shows they have been erased. However, that's only what the intruder sees; in actuality, all the files are still there.

Systems administration is handled through a GUI that is straightforward and exhaustively complete. When it's in an administrative mode, Sidewinder forces a disconnection from any outside network. This ensures that no one may slip in undetected and acquire root privileges.

If the whole process seems too daunting, SCC—like most other companies in the business—provides consulting services. The company also has software that provides Sidewinder clients with a Windows-based Internet service connection together with automated search, retrieve, and sort features. Finally, SCC offers continuing information on hacking attempts on its systems.

Trusting Your Systems

Trusted Information Systems of Glenwood, Maryland, offers a complete turn-

key firewall solution called Gauntlet Internet Firewall. This is a bastion-host system built around a Pentium-powered platform and the complete suite of TIS firewall software installed and configured to your specifications, running on a customized, secure (in government parlance, "trusted") version of BSD Unix.

Gauntlet Internet Firewall includes an integrity checker for the system itself, configurable alarms, and an audit tool that reports anomalies on a timely basis. It provides a generic interface for multiple forms of user authentication, including token-based one-time password systems.

The whole package costs $15,000, which includes setting up and configuring the system, testing, and training. For non-TIS installations, the company also offers a monthly firewall security audit, Internet gateway penetration testing, and an Internet gateway security survey.

For the do-it-yourself organization, TIS offers a Firewall Toolkit that contains source code for bare-bones versions of its logging facilities, E-mail gateway, an FTP gateway proxy, a Telnet proxy, an Internet service access-control server, and a generic pass-through plug gateway. All are available free under license from TIS on the Internet at ftp.tis.com (look for fwtk in /pub/firewalls/), or you can point your Web browser at http://www.tis.com/.

The Digital Difference

Digital Equipment (Maynard, MA) is another vendor that provides a soup-to-nuts security service called Digital's Firewall Service. The service includes consulting, software installation, hardware (if required—priced separately), training, and support. Other optional services include configuration of public-domain software, customized applications gateways, cryptographic and authentication consulting, and general-computer and network-security consulting.

The standard Firewall Service setup requires three computers, which Digital calls Gatekeeper, Gate, and Mailgate. Gatekeeper resides on the external network (what the company calls the red subnet), Mailgate resides on the network you're trying to protect (the blue subnet), and Gate resides on both. In this way, a screened subnet is established that isolates your system from the Internet or whatever public system you are trying to keep at bay. The screening software runs on Gate, a secure host. There are no user accounts (only system administration accounts) on

any of the hosts, and the applications loaded are the customized Unix utilities that Digital provides to pass acceptable packets back and forth over the link (see the figure "Digital's Three-Way Isolation").

Gatekeeper is the doorway to the outside world. It is the root DNS (Domain Name System) server of your system for the Internet, and it is where your applications gateways or proxies reside. Gatekeeper would generally be configured to accept log-ins only from trusted hosts, and the packets from these are screened according to whatever security policies you have established. Gatekeeper records all log-in attempts to the system, and it can be programmed to send an alert to the systems administrator in the case of repeated unsuccessful attempts.

If Gatekeeper is compromised, damage is limited to that single system, because Gate doesn't accept log-ins from external systems. Gatekeeper does not store the security screening policies of your network—that's one of Gate's jobs—so an intruder can't get into Gate and change them.

Besides maintaining your system's screening policies, Gate logs all attempts to connect with your internal hosts, as well as all successful log-ins to Gatekeeper and any requests for remote connections across the firewall. You can set alarm parameters to inform the systems administrator any time there is a problem. Optionally, Gatekeeper can be configured to require hand-held authentication tokens, such as Security Dynamics' (Cambridge, MA) SecurID or Digital Pathways' (Mountain View, CA) Software SecureNet Key, for successful log-in.

For E-mail, the Firewall Service uses a sendmail proxy to pass messages across the firewall. All mail between internal and external addresses is routed through Gatekeeper. If incoming mail is destined for a host on an internal network running TCP/IP, then Gatekeeper forwards the mail (through Gate) to that host. If mail is destined for a host that's not running TCP/IP, then Gatekeeper forwards it to Mailgate, which serves as a gateway to other protocols. Outgoing mail is forwarded through Mailgate and then on to Gatekeeper. Mail destined for another internal address never leaves the in-

ternal net, although it is routed through Mailgate if delivery over TCP/IP is required.

Digital's consulting services include risk assessment and impact analysis, development and implementation of policies and procedures, and security hardware and software at all levels, from the individual

Digital's Three-Way Isolation

Internet

Red subnet

Gatekeeper

Gate

Mailgate

Blue subnet

Internal network

Digital's Firewall Service uses a bastion host (Gate) to create an isolation network that sits between the external (red) and internal (blue) nets. This middle net talks to the outside world only through another secure host, Gatekeeper. It connects to the internal net either directly (for TCP/IP communications) or through a third host, Mailgate, that converts between TCP/IP and other protocols, such as DECnet.

desktop system to an entire network. In addition, they can provide training ranging from user awareness to security system management.

Pricing on Firewall Service packages starts at $20,000, which includes the operating system, applications gateways for connections to E-mail (SMTP), file transfer (FTP and Archie), remote terminal access (Telnet), client/server information services (Gopher or World Wide Web), and Notes.

A Different Approach

CheckPoint Software Technologies (Ramat Gan, Israel, and Lexington, MA) maintains that combining the packet-filtering and applications gateway approaches into a single entity is better than using either one alone. It has done just that with its $5000 CheckPoint FireWall-1 package and then topped it off with a GUI and rule-set editor that automates the process of creating the security rules for your network (for a hands-on review of FireWall-1, see page 171).

FireWall-1 has two parts: a packet-filter module that handles the actual implementation of security policy and a control module that controls and monitors one or more packet modules. Both modules may reside on the same host or on separate machines. If the modules are on different systems, communication between the two hosts is authenticated with a one-time password scheme.

While FireWall-1's design has several unique features, perhaps the most significant is its GUI. The control module used to configure FireWall-1 is an Open Look X11R5 window system, and there's a complementary set of command-line utilities for use with standard terminals. The object-oriented interface consists of the following five basic parts:

• The *Network Objects Manager* handles the definition of the various components of your network. These include, but are not limited to, the FireWall-1 hosts, servers, workstations, routers, domains, networks and subnets, databases, and more. You can also define your own objects by combining two or more standard objects.
• The *Services Manager* is the Network Objects Manager's counterpart for services. All network services are screened and controlled, whether or not you define them. The list of preloaded services is extensive, and you can define your own using simple expressions and macros. You can group services together much as you can objects.
• Once services are defined, you put them in the *Rule-Base Manager* and define your security policy rules. A high-level language makes producing your rule-set relatively easy.

Every rule has four parts: match, action, track, and target. Match specifies which

Outstanding.

Rene Magritte, Chateau des Pyrenees, ©1993 C. Herscovici/ARS, New York

Once in a while, something is created that goes above and beyond the ordinary. Something better than the rest.

In the realm of copy protection locks, the Hardlock™ copy protection system rises above the others in securing your applications against unauthorized use. Hardlock is the only lock that uses a programmable algorithm, far more complex to decode than simply reading the contents of a memory chip. Hardlock also features selectable anti-debugging and reverse engineering protection as well as protection against hardware emulators, which no other lock has.

Security. Quality. Technology.
Hardlock is state of the art.

Call us to find out more about how Hardlock can provide your masterpiece with the security it deserves.

1-800-562-2543

GLENCO
ENGINEERING INC.
SERVING THE SOFTWARE INDUSTRY SINCE 1979
Software Protection · Data Security
Phone 708-808-0300 · Fax 708-808-0313

New!
CD-Crypt
for secure
CD-ROM
software
distribution

For DOS, Windows, Windows NT, OS/2, Unix, Xenix, DES   Single User, Network, CD-ROM Applications and More

For International information circle 77 on Inquiry card.

Circle 93 on Inquiry Card.

State of the Art

communications attempts (source, destination, and service) are included in the rule. Action choices are accept, reject, or drop. Track defines what type of record, if any, should be kept of the communication. Target puts the defined rule in the appropriate packet-filter module or network location. The Rule-Base Manager's default is to drop any packet that's not explicitly defined.

• The *System Status Monitor* provides an instantaneous look at any filter module activity, including packet statistics. SNMP agent support lets this information pass to other management programs.

• The *Log Viewer* lets you view and manipulate any logged statistics, either historically or in real time.

Leave the Driving to Them

Bolt, Beranek, and Newman, the original Internet service provider based in Cambridge, Massachusetts, recently introduced what it calls the Internet Site Patrol, a turnkey package that includes consulting, hardware, software, training, and remote management services. Cost of the service starts at $1500 a month for a single gateway.

BBN's package includes an internal filtering router as backup for a two-Ethernet bastion host. The bastion host provides proxy-level services using code created by TIS. On the software side, BBN uses a mix of its own and TIS's software. A customer can make up to two software configuration changes per week free of charge, and there's 24-hour phone support for software and hardware problems. If issues cannot be resolved over the phone, BBN will have someone on-site within hours. What makes the BBN package unique is that it will manage the firewall, so the customer needn't do even that.

Different Strokes

We've looked at five different products and services that provide different approaches to firewall implementation, with varying degrees of protection. While none of this initial crop is designed (or priced) to accommodate an individual user, each can be an important component in a comprehensive, enterprise-wide program of information protection. The choice is yours. ■

John Bryan is a freelance technology writer and consultant who is based in San Jose, California. You can contact him on the Internet at 5051339@mcimail.com or on BIX c/o "editors."

BYTE
Special Report

Client/Server Computing

Today's clients and servers run on the LAN. Tomorrow's will roam the intergalactic WAN. Here's how.

Place enclosure on stand

Add a power supply…

Or two

Even optical

Snap on door

Connect to PC

INTERGALACTIC
Client/Server Computing

**ROBERT ORFALI,
DAN HARKEY, AND
JERI EDWARDS**

If you're like most of us, just now comfortable with your Ethernet LAN and local database server, you may not want to hear that the industry is poised for a second client/server revolution. Fasten your seat belts, because this next one will prove to be just as traumatic as the first, which chain-sawed mainframe applications into client and server pieces. Exponential network growth and network-aware multithreaded desktop operating systems augur a transition

> **In today's Ethernet era, there are four overlapping client/ server paradigms— SQL databases, TP monitors, groupware, and distributed objects. The next big wave is about to catapult them all into the intergalactic client/server era.**

from Ethernet client/server to intergalactic client/server, an era in which servers are plentiful instead of scarce (because every client can *be* a server) and in which proximity no longer matters.

Is the client/server infrastructure ready for intergalactic prime time? Can our NOSes (network operating systems) deal with millions of machines that can be both clients and servers? What kinds of applications will live on the global network? How will these applications be created, deployed, and managed? We'll answer these questions with a sweeping overview of today's operating systems and four dominant client/server models—SQL databases, TP (transaction processing) monitors, groupware, and distributed objects.

New Age Operating Systems

In the Ethernet era, operating systems were clearly segregated along client/server lines. Client operating systems managed the desktop; server operating systems managed shared resources. But the intergalactic era requires new hybrid operating systems that do both jobs well. The operating system must provide robust 32-bit preemptive multitasking that protects applications from one another. Clients and servers alike need threads to react quickly to events originating on the desktop and in the global network.

Clients will sport an OOUI (object-oriented user interface), which is a "place" for integrating multiple "things" that run concurrently. Things are on-screen objects that resemble their real-world counterparts. Users interact with things directly, and things exchange information by means

of drag-and-drop and live links. Technologies such as OpenDoc and OLE 2 further the OOUI paradigm, enabling users to assemble, link, script, store, and transport places and the things that they contain.

The client must also run the thousands of existing desktop applications, including DOS, Windows, OS/2, and Macintosh programs, as well as the thousands of device drivers that users have acquired. It's a daunting task, but that's what it takes to be an integrating client platform in this new era.

Of course we can't afford to ship a system administrator with every $99 operating system, so we need push-button CD-ROM installation, with dynamic discovery and configuration of resources. To achieve shrink-wrapped client/server plug and play, operating systems will bundle the required middleware. This will include protocol stacks, NOSes, resource binderies, and security features. Some of them may even have production-strength databases, TP monitors, workflow engines, and ORBs (Object Request Brokers).

To be effective as a server, a new age operating system must be upwardly scalable. It should minimally be able to exploit shared-memory SMP (symmetric multiprocessing) hardware. However,

ubiquitous WAN communication also creates the need for massively parallel, "shared-nothing" clustered servers that can service hundreds of thousands of clients and manage tons of data—video on demand, document databases, high-volume transaction processing, and information warehouses.

The New Age NOS

NOSes have always been in the business of hiding the location of resources from applications. But in the intergalactic client/server era, they must become real Houdinis, creating the illusion of a single system image across, potentially, millions of hybrid client/server machines. Here are some of the elements of that illusion:

• *Location transparency.* Users, services, and resources join and leave the network constantly, but they are never tied to fixed locations. In this continual flux, the NOS global directory brings people, programs, and things into conjunction to perform work. The global directory is a distributed, replicated object database. Distribution means that autonomous administrative domains can exist. Replication enhances availability and performance. Object orientation enables the directory to grow organically, like the real-world structures it represents.

• *Namespace transparency.* Everything on the global network must appear to belong to the same namespace. Names must resolve uniquely within a given context or

naming authority, but the NOS can grow a tree of federated namespaces, each with autonomous naming authority. You can think of the telephone system's area codes as federated namespaces.

• *Administrative transparency.* The NOS must appear to integrate with the local operating system's management services and provide replication transparency. If a naming directory is shadowed on many machines, for example, it's up to the NOS to synchronize updates. The NOS must also shield users from network failures, transparently handle retries and session reconnects, and synchronize clocks on geographically dispersed machines.

• *Secured-access transparency.* Users have to be able to access any server resource

A Simple Client/Server Model

In our book *Essential Client/Server Survival Guide*, we developed a simple model to help you navigate through the client/server maze. The model has three building blocks: a client, a server, and the slash (/) that ties the client to the server. Let's go over the pieces:

The client building block runs the client side of an application. It provides the GUI or OOUI (object-oriented user interface) element and some form of DSM (distributed system management).

The middleware building block—the slash—runs on both the client and server sides of an application. We broke this block into four categories of middleware: transport stacks, NOSes (network operating systems), DSM, and service-specific middleware. NOSes and transport stacks provide the basic communications foundation for all middleware. DSM runs on every node in a client/server network; it requires its own middleware on top of the NOS to carry messages between managing stations and managed stations. The service-specific middleware depends on the application model. Database applications use some form of SQL middleware, along with such de facto standards as ODBC (Open Database Connectivity), DRDA (Distributed Relational Database Architecture), RDA (Remote Database Access), Oracle Glue, and X/Open's CLI (call-level interface). TP (transaction processing) monitors use some form of transactional RPC (remote procedure call) or peer-to-peer middleware. Groupware applications typically use E-mail. And distributed objects typically use an ORB (Object Request Broker).

The server building block runs applications that manage shared resources. The block shows the four contending server application models: SQL databases, TP monitors, groupware, and distributed objects. The server side also contains a DSM element.

Note that the three blocks can all be running on the same machine, because any node can be both a client and a server. Server-to-server interactions are usually client/server in nature—servers are clients to other servers. However, some server-to-server interactions require specialized server middleware. For example, a two-phase commit protocol may be used to coordinate a transaction that executes on multiple servers.

The Client/Server Software Infrastructure

With four competing paradigms—SQL databases, TP monitors, groupware, and distributed objects—the middleware that connects clients to servers has grown dauntingly complex.

from anywhere, including hotel rooms, offices, homes, and cellular phones, using a single log-on. Security must be built on mutual distrust. Clients must prove to servers that they are who they claim to be and vice versa by appealing to a trusted third party. MIT's Kerberos, which is the DCE (Distributed Computing Environment) security system, works this way. After authentication, the server applications must use ACLs (access control lists) to regulate clients' access to functions and data.

• *Communications transparency.* Modern NOSes are learning to hide the complexities of multiple protocols and dissimilar data representations behind a set of abstractions for interprocess communication. All offer peer-to-peer conversational interfaces, and most provide some form of RPC (remote procedure call) that makes a server appear to be one function call away. Another model—message queuing, or MOM (message-oriented middleware)—can be incredibly helpful when clients and servers can tolerate communications delay. Current NOSes generally don't come with MOM, but it's available as an add-on.

The Ethernet-era NOSes, including NetWare 3.x, LAN Server, LAN Manager, and Windows NT Server, were built for LANs with small numbers of servers. It would be suicidal to deploy them in intergalactic environments. However, a new generation of NOSes is on the horizon, and these systems are increasingly capable of the Houdini magic needed for intergalactic environments.

The main contenders are the OSF's (Open Software Foundation) DCE, Novell's NetWare 4.x, and Sun's ONC (Open Network Computing). DCE is a technically superior NOS that meets most of the requirements; it's also the strategic NOS for major intergalactic players, including Digital Equipment, Hewlett-Packard, IBM, Microsoft, and Tandem.

Novell dominates today's client/server environments. However, the transition to NetWare 4.x has so far been a difficult one. ONC, which is entrenched on millions of Unix nodes, influenced the development of the Internet's communications infrastructure.

New Age Client/Server Applications

Everyone's imagining a planetary electronic mall, with virtual boutiques, department stores, bookstores, brokerage services, banks, travel agencies, and more. Like a Club Med, the mall will issue its

BYTE/Osborne Books are Available at the Following Locations

ALABAMA
Madison
Madison Books
and Computers
PH: 205-772-9250
FAX: 205-461-8076

ARIZONA
Tempe
Computer Library
PH: 602-820-0458
FAX: 602-491-7971

CALIFORNIA
Berkeley
Cody's Books
PH: 800-479-7744 (In CA)
800-995-1180 (Outside CA)

Cupertino
Computer Literacy
Bookshops
PH: 408-973-9955

Stacey's Professional
Bookstore
PH: 408-253-7521
FAX: 408-253-5861

Irvine
Irvine Sci-Tech Books
PH: 800-229-9514
FAX: 714-733-0122

Los Angeles
ASUCLA Students Store
PH: 800-456-0019
PH: 310-206-4041
FAX: 310-825-0382

Opamp Technical Books
PH: 800-468-4322
FAX: 213-464-0977

Palo Alto
Stacey's Professional
Bookstore
PH: 415-326-0681
FAX: 415-326-0693

Stanford Bookstore
PH: 800-673-2348
FAX: 415-322-2021

Sacramento
Tower Books
PH: 916-481-6600

San Diego
San Diego Technical Books
PH: 800-346-0071
FAX: 619-279-5088

San Francisco
Stacey's Professional
Bookstore
PH: 800-926-6511
E-mail: staceysbk@aol.com

San Jose
Computer Literacy Bookshops
PH: 408-435-1118
E-mail: info@clbooks.com

San Luis Obispo
Earthling Bookshop Inc.
PH: 805-543-7951

El Corral Bookstore
PH: 805-756-5315
FAX: 805-756-5351

Santa Barbara
Chaucers Bookstore
PH: 805-682-6787
FAX: 805-682-1129

Earthling Bookshop
PH: 805-965-0926
FAX: 805-962-6620

UCSB Bookstore
PH: 805-893-2082

Stanford
Stanford Bookstore
PH: 800-533-2670
FAX: 415-322-1936

Sunnyvale
Computer Literacy
Bookshops
PH: 408-730-9955

COLORADO
Boulder
BiblioTek
PH: 303-443-7037

Colorado Bookstore
PH: 303-442-5051
FAX: 303-444-6604

Colorado Springs
McKinzey-White
Booksellers
PH: 719-590-1700

Denver
Auraria Book Center
PH: 303-556-3230
FAX: 303-556-3736

BiblioTek
PH: 303-534-3460

Tattered Cover Bookstore
PH: 800-833-9327
FAX: 303-399-2279
E-mail:
books@tatteredcover.com

Englewood
Softpro Books
PH: 303-740-7751
FAX: 303-740-8152

Longmont
United Techbook Co.
PH: 303-651-3184
FAX: 303-651-3405

CONNECTICUT
New Haven
Yale Co-op
PH: 800-ELI-YALE
FAX: 203-772-3665

**DISTRICT
OF COLUMBIA**
Washington D.C.
Reiter's Scientific
& Professional Books
PH: 00-537-4314
FAX: 202-296-9103

GEORGIA
Atlanta
Oxford Bookstore
PH: 404-262-9975
FAX: 404-364-2729

HAWAII
Honolulu
Honolulu Book Shops
PH: 808-847-5551
FAX: 808-841-4674

University of Hawaii
Bookstore
PH: 808-956-4338
FAX: 808-956-4323

ILLINOIS
Naperville
Books & Bytes, Inc.
PH: 708-416-0102
FAX: 708-416-0375

INDIANA
Ft. Wayne
Majerek's Readers World
PH: 219-482-1531

MASSACHUSETTS
Boston
Charlesbank Bookshops
PH: 617-236-7442
FAX: 617-236-7418

Waterstone's Booksellers
PH: 617-859-7300
FAX: 617-437-0997

Burlington
Softpro
PH: 617-273-2919
FAX: 617-273-2499
E-mail:
books@softproeast.com

Cambridge
Harvard
Co-Operative Society
PH: 617-499-2000
FAX: 617-868-7038

Quantum Books
PH: 617-494-5042
FAX: 617-577-7282
E-mail:
quanbook@world.std.com

Wordsworth Books
PH: 617-498-0080
FAX: 617-354-4674

Newton Highlands
New England
Mobile Bookfair
PH: 617-527-5817
FAX: 617-527-0113

Worcester
Tatnuck
Bookseller/Databooks
PH: 800-642-6657
FAX: 508-756-9425
E-mail:
databooks@delphi.com

MICHIGAN
East Lansing
Grand River Books
PH: 517-333-3666

Flint
Young & Welshans
Book Sellers
PH: 810-732-0626
FAX: 810-732-2393

MINNESOTA
Minneapolis
Baxters Books
PH: 612-339-4922
PH: 800-626-1049
FAX: 612-339-6134

NEW HAMPSHIRE
Hanover
Dartmouth Bookstore
PH: 800-675-3616 (in NH)
800-624-8800 (outside NH)
FAX: 603-643-5170

NEW MEXICO
Albuquerque
Page One, Inc.
PH: 800-521-4122
FAX: 505-294-5576

NEW YORK
New York
Barnes & Noble
Fifth Ave.
PH: 212-807-0099

Computer Book Works
Warren St.
PH: 212-385-1616
FAX: 212-385-8193

McGraw-Hill Book Store
Ave. of the Americas
PH: 212-512-4100
FAX: 212-512-4105

OHIO
Cincinnati
University of
Cincinnati Bookstore
PH: 513-556-1800
FAX: 513-556-5555

OREGON
Portland
Powell's Technical Books
PH: 503-228-3906
FAX: 503-228-0505

PENNSYLVANIA
Pittsburgh
Book Center -
U. of Pittsburgh
PH: 412-648-2321
FAX: 412-648-1902

RHODE ISLAND
Providence
Brown Bookstore
PH: 401-863-3168
FAX: 401-863-2233

TEXAS
Dallas
Taylors Bookstore Ltd.
PH: 214-357-1700
FAX: 214-956-8012

VIRGINIA
Vienna
Computer Literacy
Bookshops
PH: 703-734-7771
E-mail: sales@tc.clbooks.com

WASHINGTON
Bellevue
Tower Books
PH: 206-451-1110
FAX: 206-454-0453

University Bookstore
PH: 206-646-3300
FAX: 206-634-0810

Seattle
University Bookstore
PH: 206-634-3400
FAX: 206-646-3340

WISCONSIN
Madison
University Bookstore
PH: 608-257-3784
PH: 800-993-2665
FAX: 608-257-9479

Milwaukee
Harry W. Schwartz
Bookshops
PH: 414-274-6460
PH: 800-236-7323
FAX: 414-274-6408

own electronic currency to facilitate round-the-clock shopping and for business-to-business transactions. Electronic agents will roam the network, looking for bargains and negotiating with other agents. Billions of transactions and oceans of multimedia data will flow through the network every day.

Clearly, we're not talking about today's Internet, where users surf hypertext webs of HTML-tagged (Hypertext Markup Language) information. The volume and complexity of transactions, and the richness of the data on which they operate, will create the need for new enabling technologies, including:

• *Rich transaction processing.* We'll need nested transactions that can span servers, transactions that execute over long periods of time as they travel from server to server, queued transactions for secure business-to-business dealings, and "sagas" that can chain together many pieces of work and selectively undo some of the effects of a transaction. Most nodes on the network should be able to participate in a secured transaction. Superserver nodes will handle the massive transaction loads.
• *Roaming agents.* Consumers will have personal agents that look after their interests. Businesses will deploy agents to sell their wares on the network. Sniffer agents will look for trends and gather statistics. Agent technologies include cross-platform scripting engines, work-flow engines, and an infrastructure that lets agents live on any machine on the network.
• *Rich data management.* From anywhere on the network, we'll create, store, view, and edit compound documents with multimedia content. Most nodes will offer compound document technology (e.g., OLE or OpenDoc) for local document management. Superservers will provide repositories for storing and distributing massive numbers of documents. Of course, we can't forget about structured data (e.g., SQL databases) either.

The Four Client/Server Application Models

What technology base will be used to create these intergalactic client/server applications? The four competing paradigms are SQL databases, TP monitors, groupware, and distributed objects. Each one of them can create complete client/server applications, each provides tools to do that (some more than others), and each introduces its favorite form of middleware. In the remainder of the article, we'll explain what each contender does well and fearlessly pick a winner.

SQL DATABASES

SQL database servers dominate the client/server landscape today. SQL began as a declarative language for manipulating data using 10 simple commands. But as SQL applications moved to more demanding client/server environments, it became clear that just managing data wasn't enough. There was also a need to manage the functions that manipulated the data. Stored procedures, sometimes called "TP lite," met the need.

A stored procedure is a named collection of SQL statements and procedural logic that is compiled, verified, and stored in a server database. Sybase pioneered the concept of stored procedures. Now virtually all SQL vendors support stored procedures along with other SQL extensions (e.g., triggers and rules). The extensions are used to enforce data integrity, perform system maintenance, and implement the server side of an application's logic. No two vendor implementations are alike. Note that stored procedures offer minimal transaction support.

Because SQL standards seem to lag vendor implementations by at least five years, almost everything that's interesting in client/server database technology is nonstandard. This includes database administration, data replication, stored procedures, user-defined data types, client APIs, and the formats and protocols on the network. Thus, the best you can do in heterogeneous database environments is to create a loose federation of databases whose least common denominator is typically dynamic SQL.

In defense of SQL, we can say that it is easy to create client/server applications in single-vendor/single-server environments. A wealth of GUI tools makes SQL applications easy to build, and SQL is familiar to millions of programmers and users.

TP MONITORS

You can't create mission-critical applications without managing the programs (or processes) that operate on data. That's why, in the mainframe world, a TP monitor comes with every mission-critical database. TP monitors manage processes and orchestrate programs by breaking complex applications into pieces of code called *transactions.* The modern client/server incarnations of TP monitors haven't dominated the Ethernet era, but they'll be major players in the intergalactic era. It's not far-

fetched to assume that every machine on the network will have a TP monitor to represent it in global transactions.

Transactions are more than just business events; they've become an applications design philosophy that guarantees robustness in distributed systems. Under the control of a TP monitor, a transaction can be managed from its point of origin—typically on a client—across one or more servers and back to the originating client. When a transaction ends, all parties involved agree that it either succeeded or failed.

The transaction is the contract that binds the client to one or more servers. It's the fundamental unit of recovery, consistency, and concurrency in a client/server system. Of course, all participating programs must adhere to the transactional discipline; otherwise, a single faulty program can corrupt an entire system. In an ideal world, all client/server programs will be written as transactions.

Transaction models define when a transaction starts, when it ends, and what the appropriate units of recovery are in case of failure. The flat-transaction model has long been the workhorse of the current generation of TP monitors (and other transactional systems). In a flat transaction, all work done within a transaction's boundaries occurs at the same level. The transaction starts with a `begin_transaction` and ends with either a `commit_transaction` or an `abort_transaction`. It's all or nothing—there's no way to commit or abort parts of a flat transaction. However, newer transaction models offer finer control of a transaction's threads and can more closely mirror their real-world counterparts.

Most alternatives to the flat transaction extend the flow of control beyond the simple unit of work, either by chaining units of work in linear sequences of minitransactions or via nested subtransactions (see the figure "Nested Transactions"). A subtransaction's effects become permanent after it issues a local commit and all its ancestors commit. If a parent transaction aborts, all descendant transactions abort whether or not they issue local commits. The beauty of it is that subtransactions can run on different nodes.

TP monitors were invented to run applications that serve thousands of clients. By interjecting themselves between clients and servers, TP monitors can manage transactions, route them across systems, load-balance their execution, and restart them after failures. A TP monitor can manage transactional resources on a single

1. Standardize On A Common GUI Interface That Makes It Easier To Use, Reduce Training Costs And Boost Productivity.

2. Protect All Your Data With Bulletproof Security Software That Supports Every Platform And Operating System Across Your Entire Enterprise.

ERROR
SECURITY VIOLATION

UNIX

AS/400

Data Center Administration

Automated Production Control

MVS

CA-UNICENTER

OS/2

Automated Storage Management

Security, Control And Audit

WINDOWS NT

NETWARE

Performance Management And Accounting

3. Choose A Client/Server Application That's Interoperable, Scalable and Portable.

4. Demand An Open Solution That Supports All Key Industry Standards And Protocols.

5. Partner With The World's Leading Software Company.

6. Leverage Your Systems Management Experience, Expertise And Personnel.

7. Replace Disparate Point Products With A Single, Integrated, Comprehensive Solution.

Presenting The 7 Commandments Of Systems Management Software.

CA World '95 July 16-21, 1995 New Orleans. Call 1-800-CA-INFO95 for information.

CA-Unicenter

When choosing your systems management software, you shouldn't have to make a leap of faith. And with new CA-Unicenter, you don't have to. Because CA-Unicenter is the single, integrated systems management solution that solves all of the challenges of managing client/server computing across heterogeneous platforms.

So call today for more information about the systems management software you can believe in: CA-Unicenter.

Call 1-800-225-5224, Dept. 10500 For More Information And To Register For The CA-Unicenter Software Giveaway.

CA-Unicenter®
Integrated Client/Server Systems Management Software

COMPUTER ASSOCIATES
Software superior by design.

© Computer Associates International, Inc., Islandia, NY 11788-7000. Offer good in U.S. and Canada only. All other product names referenced herein are trademarks of their respective companies. Requirements for Software Giveaway available through toll-free number.

Circle 289 on Inquiry Card.

IF THIS IS HOW YOU

◉ Datapro on CD-ROM

Now you gain virtually instant access to Datapro's product and technology services from your own PC, with Datapro on CD-ROM.

Menu-driven software makes it simple to browse through the table of contents or search by keywords and subject to get fast answers to all your technology questions. And you can simultaneously view multiple reports and cut and paste between applications to custom-tailor analyses. Available in both Windows and DOS.

Nested Transactions

Nested transactions can capture the true complexity of real-world business processes while distributing the transaction load across many server nodes.

server or across multiple servers, and it can cooperate with other TP monitors in federated arrangements. TP monitors also perform a "great funneling act" that helps the operating system and server resource managers deal with large numbers of clients (see "Scale Up with TP Monitors" on page 123).

X/Open and the OMG (Object Management Group) have created complementary standards that define how TP monitors interact with applications, resource managers, and other TP monitors in both procedural and distributed-object en-

vironments. Some good tools are available to help you create TP-monitor applications.

TP monitors are probably overkill in single-server, single-vendor departmental applications. That's one reason they've been slow to take off. Moreover, vendors haven't come to grips yet with the realities of the shrink-wrapped software market, and they haven't been able to explain the advantages TP monitors offer. The intergalactic era will make those advantages increasingly self-evident.

GROUPWARE

Groupware comprises five foundation technologies geared to support collaborative work: multimedia document management, work flow, E-mail, conferencing, and scheduling. Groupware isn't another downsized mainframe technology. It's a genuinely new model of client/server computing. It helps users collect unstructured data—including text, images, faxes, mail, and on-line conference proceedings—and organize that data as a set of documents. It

For 1200 dpi printing,
there's really only
one choice.
There are, however,
several options.

The new Optra family
from Lexmark.

At Lexmark, we don't limit top-of-the-line features to our top-of-the-line printers.

The new Optra™ family is the first full line of network-ready, desktop laser printers capable of true 1200 x 1200 dpi* printing. Each Optra laser printer, from the extremely affordable 12 ppm Optra R to the extremely powerful 16 ppm Optra Lxi, delivers four times the sharpness of ordinary 600 dpi printers.

No matter which option you choose, our break-through MarkVision™ printer utility is standard. Users and LAN managers can control and configure every Optra on the network, and monitor job statistics, right from Windows™ or Macintosh workstations.

PostScript™ Level 2 and enhanced PCL° 5 emulations are also standard. Which means that every Optra prints incredibly sharp graphics and beautiful, crisp text from virtually any of your applications.

All this value and innovation from Lexmark, a former division of IBM. For more information, call us at **1 800 891-0399, ext. 230.**

Thanks to the new Optra family, you can choose the best printer, rather than the best compromise.

LEXMARK™

ADVANCING THE ART OF PRINTING

further enables users to view those documents, store them, replicate them, and route them anywhere on the network. The multimedia document is to groupware what a table is to a SQL database—the basic unit of management.

Groupware excels in the art of document database management and makes effective use of E-mail, which is its preferred form of middleware. E-mail is one of the easiest ways for electronic processes to communicate with humans. Asynchronous by nature, it's a good match for the way businesses really work. E-mail is ubiquitous, with over 50 million globally interconnected electronic mailboxes.

To appreciate what makes groupware technology so different, consider Lotus Notes, the premier groupware product in the industry (see "Document Repositories" on page 131). The secret of Notes' success is that its whole is much more than the sum of its parts. Like all good groupware, Notes makes effective use of E-mail and can manage document databases in a client/server fashion.

What makes Notes revolutionary, however, is widespread replication of these databases. What happened to locking and data integrity, the watchwords of SQL databases? Notes doesn't care—it was more important to get the information out. Now, that's revolutionary! With release 3, Notes introduced a level of version control that provides adequate protection for document-oriented applications. But Notes is not recommended if you require concurrency or immediate updates.

Managing business processes by means of work flow is another revolutionary aspect of groupware. In a work flow, data (and sometimes functions) passes from one program to the next in structured or unstructured client/server environments. Modern work-flow software electronically simulates real-world collaborative activity. Work can be routed in ways that correspond to interoffice communications. You can create sequential routes, parallel routes (i.e., alternate paths), routes with feedback loops, circular routes, and more.

A good work-flow package lets you specify acceptance criteria for moving work from one stage to the next. So, work flow brings the information to the people (and programs) who can act on it. It can also coordinate existing software and track processes to make sure the work gets done by the right people.

Groupware provides many of the components we need for creating intergalactic client/server applications. The technology is also starting to encroach on its competitors' turf. For example, Lotus Notes—via DataLens—can access information that's stored in SQL databases; it also gives SQL applications access to Notes data by way of ODBC (Open Database Connectivity).

New tools (e.g., Notes VIP) integrate GUI-building client facilities with data that can be accessed from document or SQL databases. At its best, groupware can flexibly combine different client/server technologies and adapt to the way people do business in both structured and ad hoc settings.

DISTRIBUTED OBJECTS

Distributed-object technology promises the most flexible client/server systems. This is because it encapsulates data and business logic

KFC offers a full line of 9",14", 15",17",20" and 21" green monitors

Now we present the innovative GREEN MONITORS to protect our environment with you. These monitors are compliant with VESA DPMS standard to manage the power system itself. It could save 95% power consumption when your computer is not in operation.

For optimizing the display quality, the front panel control allows you to adjust the pincushion distortion, keystone effect, image rotation and full screen display.

Over ten years of design and manufacturing experience ensure that KFC monitors consistently meet your needs.

◆ KFC ®   ISO 9002
CERTIFICATED MANUFACTURER

HEADQUARTERS Taiwan
TEL:886-2-7548353
FAX:886-2-7548498

CALIFORNIA USA
TEL:818-3386888
FAX:818-3382080

TEXAS, USA(OEM)
TEL:512-8356776
FAX:512-8356779

JANUARY 1994
BYTE
BEST VALUE
GENERAL BUSINESS
COLOR MONITOR

energy

"The energy star emblera does not represent EPA endorsement of any product or servicer."

Circle 518 on Inquiry Card.

in objects that can roam anywhere on networks, run on different platforms, talk to legacy applications by way of object wrappers, and manage themselves and the resources they control.

When it comes to standards, distributed-object technology is ahead of the other client/server approaches. Since 1989, a consortium of object vendors—called the OMG—has been busily specifying the architecture for an open software bus on which object components written by different vendors can interoperate across networks and operating systems. The OMG currently boasts 440 member companies, and the object bus is well on its way to becoming the mother of all client/server middleware.

The secret to the OMG's success is that it defined how to specify an interface between a component and the object bus—using working technologies as a model—but didn't prescribe how to implement those specifications. Specifications are written in IDL (interface definition language), independent of any programming language. Components specify in IDL the types of services that they provide, including the methods they export and their parameters, attributes, error handlers, and inheritance relationships with other components.

IDL becomes the contract that binds clients to server components. The beauty of IDL is that it can easily be used to encapsulate existing applications. You don't have to rewrite your entire inventory of applications to take advantage of distributed-object technology.

The object bus provides an ORB that lets clients invoke methods on remote objects either statically or dynamically. If a component interface is already defined, you can bind your program to an IDL-generated stub to call its methods. Otherwise, you can discover how the interface works at run time by consulting an OMG-specified interface repository.

In late 1994, the OMG approved a set of specifications that's called CORBA 2.0 (Common Object Request Broker Architecture), which defines a TCP/IP-based ORB-to-ORB backbone. CORBA 2.0 also specifies an optional inter-ORB commu-

nications service based on DCE. The new interface-repository specification defines extensions that let the components generate universal global IDs for their interfaces—using IDL, of course—to ensure that they are unique at the intergalactic level.

In addition to defining the object bus, the OMG has specified an extensive set of ORB-related services for creating and deleting objects, accessing them by name, storing them, externalizing their states, and defining complex relationships among them. In late 1994, the OMG also defined a comprehensive set of services for transactional objects. The idea is that you'll be able to create an ordinary object and make it transactional, lockable, and persistent by having it inherit the appropriate services using simple IDL entries.

The OMG created important alliances to make sure that its standards are universally accepted. The CORBA-defined persistence service is closely aligned with the new ODMG-93 (Object Database Management Group) specifications for object databases. The object-transaction services can interoperate with X/Open-defined procedural transactions. The OMG is also working with X/Open to help define ORB-based system management interfaces, including security.

In addition, CI Labs—which is the consortium of companies responsible for OpenDoc—picked CORBA as its object model for intercomponent communications. Both Taligent and OpenStep are providing CORBA gateways for their external object communications. Even Microsoft approached the OMG in late 1994 with a proposal for an official OLE-to-CORBA gateway.

It looks like the object community may be well on the way to building an object infrastructure that can meet the demands of the intergalactic client/server era. Distributed objects with the proper component packaging and infrastructure may provide the ultimate building blocks for creating client/server solutions, including suites of co-operating business objects.

In all honesty, though, the current generation of CORBA 1.2–compliant ORBs is not ready for intergalactic prime time. We believe that CORBA 2.0 and the new object services—including transactions, locking, life cycle, naming, and persistence—must be implemented in commercial ORBs before the technology can take off. The good news is that in this case, unlike the SQL world, the standards lead rather than lag the commercial offerings.

Once distributed-object technology takes off, it can subsume all other forms of client/server computing, including TP monitors, SQL databases, and groupware. Distributed objects can do it all and do it better. Objects can help us break large monolithic applications into more manageable components that coexist on the intergalactic bus. They are also our only hope for managing the millions of software entities that will live on intergalactic networks.

As shown in the figure "Three Waves of Client/Server Computing," the Ethernet era of client/server saw a file-centered applications wave followed by a database-centered wave. TP monitors and groupware generated minor ripples. Distributed objects are the next big wave, and we believe they're the only way to make the intergalactic client/server vision real. Are you ready to catch the wave? ■

Robert Orfali, Dan Harkey, and Jeri Edwards are the authors of Essential Client/Server Survival Guide *(Van Nostrand Reinhold, 1994). Edwards is the director of transaction processing and client/server software development at Tandem Computers. Orfali and Harkey work on the application of distributed-object technology at IBM; they have worked on applied client/server technology for the last eight years. You can reach them on BIX c/o "editors."*

Three Waves of Client/Server Computing

In the Ethernet client/server era, which is still very much in progress, SQL databases, TP (transaction processing) monitors, and groupware have begun to displace file servers. But the intergalactic client/server era, with its distributed-object infrastructure, is fast approaching.

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today.
For quickest response, fax to 1-416-620-9790!

Product Category Information

To receive information for an entire category, circle the appropriate number on the adjacent card.

Hardware

Accessories/Supplies	1
Add-in Boards	2
Bar Coding	3
Communications/Networking	4
Computer Systems	5
Data Acquisition	6
Diagnostic Equipment	53
Disks & Optical Drives	7
Diskettes/Duplicators	8
Fax Boards/Machines	9
Graphics Tablets/Mice/ Pen Input	10
Keyboards	11
LAN Hardware	12
Laptops & Notebooks	13
Mail Order	14
Memory/Chips/Upgrades	15
Miscellaneous Hardware	16
Modems/Multiplexors	17
Monitors & Terminals	18
Multimedia/CD-ROM	19
PCMCIA	57
Printers/Plotters	20
Programmable Hardware	21
RAID Drive Arrays	56
Scanners/OCR/Digitizers	22
Security	52
Tape Drives	23
UPS/Power Management	24
Voice Technology	55

Software

Business	25
CAD/CAM	26
Communications/Networking	27
Data Acquisition	28
Database	29
Educational	30
Engineering/Scientific	31
Entertainment	32
Graphics	33
Macintosh	34
Mail Order	35
Mathematical/Statistical	36
Miscellaneous Software	37
On-Line Services	38
Operating Systems	39
OS/2	54
Programming Languages/ Tools	40
Security	41
Shareware	42
Software Duplication	43
Spreadsheets	44
Unix	45
Utilities	46
Windows	47
Word Processing	48

General

Books/Publications	49
Recruitment	50
Miscellaneous	51

Fill out this coupon carefully. *Please Print.*

Name _____ Title _____
Company _____ Address _____
City _____ Country _____ Code _____
Phone _____ Fax _____

A. My job function is best classified as (check one):
1 [ ] Company Management
2 [ ] IS/MIS/IT Management
3 [ ] Systems Engineering/Integration
4 [ ] Systems/Networking Consulting
5 [ ] Departmental Management (non-IS/MIS)
6 [ ] Technical Services Support
7 [ ] Other (please describe)

B. What is your organization's primary business activity at this location (check one):
8 [ ] Business Services (Finance, Banking, Insurance, Healthcare, Professional)
9 [ ] Commerce/Industry (Retail, Wholesale, Construction, Mining, Manufacturing, Transportation)
10 [ ] Reseller/OEM (VAR, VAD, Systems/Network Integrator, Computer Product Manufacturer)
11 [ ] Government (Federal, State, Municipal, Military)
12 [ ] Computer Services (Support, Training, Consulting)

13 [ ] Other (please specify): _____

C. Please indicate which specific fields of interest are important to you so that BYTE may serve you with supplemental information that best meets your needs (select all that apply):
14 [ ] UNIX and Workstations
15 [ ] Networking (LANs, WANs, and Telecommunications)
16 [ ] Multimedia
17 [ ] Reselling/Systems Integration
18 [ ] Applications Development
19 [ ] Other (explain): _____

D. My responsibilities require that I be involved with the following operating system environments (check all that apply):
20 [ ] DOS
21 [ ] OS/2
22 [ ] Mac/OS

23 [ ] UNIX (any, including Solaris)
24 [ ] Windows
25 [ ] Windows/NT
26 [ ] NetWare
27 [ ] Other (please describe):

E. The number of employees at my location and company-wide area (check one in each column):

	At My Location	Company-Wide
10,000 or more	28 [ ]	33 [ ]
5,000 to 9,999	29 [ ]	34 [ ]
1,000 to 4,999	30 [ ]	35 [ ]
100 to 999	31 [ ]	36 [ ]
Under 100	32 [ ]	37 [ ]

April 1995
Valid until June 30, 1995

90 97 98

Inquiry Numbers 1–187

```
1    2    3    4    5    6    7    8    9   10   11   12   13   14   15   16   17
18   19   20   21   22   23   24   25   26   27   28   29   30   31   32   33   34
35   36   37   38   39   40   41   42   43   44   45   46   47   48   49   50   51
52   53   54   55   56   57   58   59   60   61   62   63   64   65   66   67   68
69   70   71   72   73   74   75   76   77   78   79   80   81   82   83   84   85
86   87   88   89   90   91   92   93   94   95   96   97   98   99  100  101  102
103  104  105  106  107  108  109  110  111  112  113  114  115  116  117  118  119
120  121  122  123  124  125  126  127  128  129  130  131  132  133  134  135  136
137  138  139  140  141  142  143  144  145  146  147  148  149  150  151  152  153
154  155  156  157  158  159  160  161  162  163  164  165  166  167  168  169  170
171  172  173  174  175  176  177  178  179  180  181  182  183  184  185  186  187
```

Inquiry Numbers 188–374

```
188  189  190  191  192  193  194  195  196  197  198  199  200  201  202  203  204
205  206  207  208  209  210  211  212  213  214  215  216  217  218  219  220  221
222  223  224  225  226  227  228  229  230  231  232  233  234  235  236  237  238
239  240  241  242  243  244  245  246  247  248  249  250  251  252  253  254  255
256  257  258  259  260  261  262  263  264  265  266  267  268  269  270  271  272
273  274  275  276  277  278  279  280  281  282  283  284  285  286  287  288  289
290  291  292  293  294  295  296  297  298  299  300  301  302  303  304  305  306
307  308  309  310  311  312  313  314  315  316  317  318  319  320  321  322  323
324  325  326  327  328  329  330  331  332  333  334  335  336  337  338  339  340
341  342  343  344  345  346  347  348  349  350  351  352  353  354  355  356  357
358  359  360  361  362  363  364  365  366  367  368  369  370  371  372  373  374
```

Inquiry Numbers 375–561

```
375  376  377  378  379  380  381  382  383  384  385  386  387  388  389  390  391
392  393  394  395  396  397  398  399  400  401  402  403  404  405  406  407  408
409  410  411  412  413  414  415  416  417  418  419  420  421  422  423  424  425
426  427  428  429  430  431  432  433  434  435  436  437  438  439  440  441  442
443  444  445  446  447  448  449  450  451  452  453  454  455  456  457  458  459
460  461  462  463  464  465  466  467  468  469  470  471  472  473  474  475  476
477  478  479  480  481  482  483  484  485  486  487  488  489  490  491  492  493
494  495  496  497  498  499  500  501  502  503  504  505  506  507  508  509  510
511  512  513  514  515  516  517  518  519  520  521  522  523  524  525  526  527
528  529  530  531  532  533  534  535  536  537  538  539  540  541  542  543  544
545  546  547  548  549  550  551  552  553  554  555  556  557  558  559  560  561
```

Inquiry Numbers 562–748

```
562  563  564  565  566  567  568  569  570  571  572  573  574  575  576  577  578
579  580  581  582  583  584  585  586  587  588  589  590  591  592  593  594  595
596  597  598  599  600  601  602  603  604  605  606  607  608  609  610  611  612
613  614  615  616  617  618  619  620  621  622  623  624  625  626  627  628  629
630  631  632  633  634  635  636  637  638  639  640  641  642  643  644  645  646
647  648  649  650  651  652  653  654  655  656  657  658  659  660  661  662  663
664  665  666  667  668  669  670  671  672  673  674  675  676  677  678  679  680
681  682  683  684  685  686  687  688  689  690  691  692  693  694  695  696  697
698  699  700  701  702  703  704  705  706  707  708  709  710  711  712  713  714
715  716  717  718  719  720  721  722  723  724  725  726  727  728  729  730  731
732  733  734  735  736  737  738  739  740  741  742  743  744  745  746  747  748
```

Inquiry Numbers 749–935

```
749  750  751  752  753  754  755  756  757  758  759  760  761  762  763  764  765
766  767  768  769  770  771  772  773  774  775  776  777  778  779  780  781  782
783  784  785  786  787  788  789  790  791  792  793  794  795  796  797  798  799
800  801  802  803  804  805  806  807  808  809  810  811  812  813  814  815  816
817  818  819  820  821  822  823  824  825  826  827  828  829  830  831  832  833
834  835  836  837  838  839  840  841  842  843  844  845  846  847  848  849  850
851  852  853  854  855  856  857  858  859  860  861  862  863  864  865  866  867
868  869  870  871  872  873  874  875  876  877  878  879  880  881  882  883  884
885  886  887  888  889  890  891  892  893  894  895  896  897  898  899  900  901
902  903  904  905  906  907  908  909  910  911  912  913  914  915  916  917  918
919  920  921  922  923  924  925  926  927  928  929  930  931  932  933  934  935
```

Inquiry Numbers 936–1122

```
936  937  938  939  940  941  942  943  944  945  946  947  948  949  950  951  952
953  954  955  956  957  958  959  960  961  962  963  964  965  966  967  968  969
970  971  972  973  974  975  976  977  978  979  980  981  982  983  984  985  986
987  988  989  990  991  992  993  994  995  996  997  998  999 1000 1001 1002 1003
1004 1005 1006 1007 1008 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1020
1021 1022 1023 1024 1025 1026 1027 1028 1029 1030 1031 1032 1033 1034 1035 1036 1037
1038 1039 1040 1041 1042 1043 1044 1045 1046 1047 1048 1049 1050 1051 1052 1053 1054
1055 1056 1057 1058 1059 1060 1061 1062 1063 1064 1065 1066 1067 1068 1069 1070 1071
1072 1073 1074 1075 1076 1077 1078 1079 1080 1081 1082 1083 1084 1085 1086 1087 1088
1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105
1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122
```

Inquiry Numbers 1123–1309

```
1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139
1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150 1151 1152 1153 1154 1155 1156
1157 1158 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1170 1171 1172 1173
1174 1175 1176 1177 1178 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1190
1191 1192 1193 1194 1195 1196 1197 1198 1199 1200 1201 1202 1203 1204 1205 1206 1207
1208 1209 1210 1211 1212 1213 1214 1215 1216 1217 1218 1219 1220 1221 1222 1223 1224
1225 1226 1227 1228 1229 1230 1231 1232 1233 1234 1235 1236 1237 1238 1239 1240 1241
1242 1243 1244 1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258
1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275
1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292
1293 1294 1295 1296 1297 1298 1299 1300 1301 1302 1303 1304 1305 1306 1307 1308 1309
```

Inquiry Numbers 1310–1489

```
1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1320 1321 1322 1323 1324 1325 1326
1327 1328 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1340 1341 1342 1343
1344 1345 1346 1347 1348 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360
1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377
1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394
1395 1396 1397 1398 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411
1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428
1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445
1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462
1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479
1480 1481 1482 1483 1484 1485 1486 1487 1488 1489
```

BYTE

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today.
For quickest response, fax to 1-416-620-9790!

BYTE

TNT MAILFAST
FRG/BOS/700252/MO1
P.O. Box 66
Hounslow
TW5 9RT
England

1. For FREE product information from individual companies, circle the corresponding inquiry numbers on the Response Card!

2. Print Your Name and Address
Answer questions "A" through "E" and mail or fax card to 1-416-620-9790

3. Product information will be rushed to you from the selected companies!

Place Postage Here

BYTE

Scale Up with TP Monitors

**JIM GRAY AND
JERI EDWARDS**

I n a simple client/server system, many clients issue requests and one server responds. But when you scale up from 50 clients to 500 or more, this model breaks most operating systems. TP (transaction processing) monitors solve this scale-up problem by modifying the simple request-response flow using techniques that have evolved over the last 30 years. They also introduce tools for designing, configuring, managing, and operating client/server systems. We'll first focus on how TP monitors tackle the fundamental upsizing problem and then look at some of their configuration and management capabilities.

Dedicating one server process per client is the obvious way to build a client/server application. In this model, client applications run concurrently; if one malfunctions, others are not affected. But the server-process-per-client design has two severe scalability problems. First, there's the percentage problem. At 10 clients, each client gets 10 percent of the server; however, at 100 clients, each gets just 1 percent of the server. The shared server and its data become precious resources as the number of clients rises above 30 or so.

Then there's the polynomial explosion problem (see the figure "Process per Client: X×Y×Z Connections"). Each client typically wants to open several applications. Each application wants to open several files. So X clients opening Y applications each with Z open files results in X×Y processes and X×Y×Z connections and open files. Before long, this adds up to thousands of processes, tens of thousands of connections, and a crashed operating system.

The obvious solution to these problems was to go from a server process per client to a server process per server (see the figure "Process per Server: X+Z Connections"). Many companies implemented an efficient operating system within the operating system—folding the entire application within one operating-system process. This TP operating system implemented a private thread library and built a private file system based on the host operating system's raw file interface. Typical of this approach was IBM's original CICS, which could support hundreds of clients on a half-MIPS, 100-KB server running DOS/360. The idea was reincarnated in the 1980s by Novell's NetWare and by Sybase's SQL Server.

Sybase multithreaded a single Unix process and added stored procedures, thus getting a three times speedup and

a 10 times scale-up advantage over the process-per-client servers from Oracle, Ingres, and Informix. Today, database vendors have copied the Sybase design and offer multithreaded servers with stored procedures. Novell's NetWare file server, introduced in 1982, quickly evolved into a database and applications server. By using inexpensive threads and emphasizing the performance of the simple requests, NetWare was able to support many clients with a relatively modest server.

The process-per-server design addresses the percentage problem by offering efficient services. Novell is justly proud of NetWare's ability to service a client's disk request in under a thousand instructions, which is 10 to 100 times better than such general-purpose operating systems as OS/2, Unix, and Windows NT. The process-per-server design solves the polynomial explosion problem by having only one server process. There are only X client connections to the server; Y applications at the server collectively open only Y×Z files. These are manageable numbers.

Applications Server Partitions

CICS, NetWare, and Sybase have been extraordinarily successful, but there are problems with the process-per-

> **When client/server applications grow to accommodate hundreds or thousands of clients, today's operating systems break down. Here's how TP monitors come to the rescue.**

STEVE KELLER © 1995

Process per Client: $X \times Y \times Z$ Connections

Clients   Servers   Files

Each client can use many server processes, and each server process can use many files. The resulting polynomial explosion can rapidly overwhelm an operating system.

Process per Server: $X+Z$ Connections

Clients   Server   Files

Using threads, a single-server process drastically reduces the number of connections and open files required to support a given number of clients.

Process per Partition: $X+1$ Connections per Server

Clients   Server partitions   Files

With partitioned data, the application can scale across multiple servers. But it's tricky to partition the data, and clients have to maintain connections to multiple servers.

server design. It does not scale to shared-memory symmetric multiprocessors because the single operating-system process uses a single processor. Other processors just sit idle. If that single operating-system process faults or waits in any way, the entire server stalls. Even worse, the process-per-server design does not scale to clusters of servers.

Beyond these scalability problems, the process-per-server model has a manageability problem. The design creates a monolithic process that collapses all applications into one address space. A bug in any application can crash the server. Changing any application can impact all others.

These scaling and management problems obviously suggest the idea of a process-per-application-server partition (see the figure "Process per Partition: $X+1$ Connections per Server"). The idea is to specialize a process or processor to service a particular application function. You scale the system by adding servers for each application. If an application saturates a single server, you partition the application data and dedicate a server to each partition.

The process-per-application-server-partition technique is widely used to scale up CICS, NetWare, Sybase, and Oracle applications. The difficulty is that it reintroduces the polynomial explosion problem. The clients must connect to each application-server partition, log on to it, and main-

tain a connection with it. The client code needs to route requests to the appropriate partition.

It is not easy to partition most applications. A particular request may touch many partitions. There are often central files or resources that all partitions or applications (e.g., the customer list, the price list, and the bindery) use. Partitioning such resources is not possible, so they must be replicated or managed by a shared server. Nonetheless, process-per-application partition is the most widely used scalability technique today.

All the solutions described so far involve two kinds of processes: clients or servers. These are generically called two-ball designs. All the two-ball designs expect the client to find the servers and route requests to the appropriate server. Each server authenticates the client and manages the connection to the client.

Routers: A More Scalable Design

The three-ball model introduces a router function (see the figure "Three-Ball Model: Routers Have $X+A$ Connections"). The client connects to a router, and the router brokers client requests to servers. The client is authenticated once and sends all its requests through a single connection to its router. This design scales by adding more routers as the number of clients grows.

Routers typically create and manage pools of application-server processes. All members of a process pool provide identical services. A pool can be distributed across the several nodes of a cluster; the routers balance the load. Each application can have a separate server pool. The router can run different pools (applications) at different priorities to optimize response time for simple requests. If a server fails, the router redirects the request to another member of the pool. This arrangement provides load-balancing and transparent server fail-over for clients.

IBM's IMS, built in 1970, was the first three-ball system. It had a single router process. With time, Tandem (Pathway, 1979), Digital Equipment (ACMS, 1981 and RTR, 1987), AT&T (Tuxedo, 1985 and Topend, 1991), and Transarc (Encina, 1993) generalized the ideas to provide many additional features.

The process-per-client model had the virtue of implementation simplicity, and each client benefitted by having its own server process. However, the design did not scale up because of the percentage problem and the polynomial explosion problem. The two-ball model collapsed all the applications together, thereby solving

these two problems but creating other scalability problems.

The three-ball model multiplexes the several clients down to a few server processes. This solves the polynomial explosion problem, but the percentage problem remains an issue. The three-ball model uses the operating system to provide server processes. The benefit is that the router and applications can use symmetric multiprocessors and clusters. They can scale up to large systems with high throughput. If the router is designed carefully and the operating system dispatches well, the three-ball model can compete with the uniprocessor two-ball systems.

The three-ball model allows the applications designer to use either a process per server CPU, a process per application, or a process per client, which is the most interesting case. By dynamically connecting the client to a server on an as-needed basis, the three-ball router increases the duty cycle on each server. This solves the polynomial explosion problem while still allowing the application to have simple interactions with the client.

Initially, classic CICS (on the mainframes) was implemented as a two-ball model. When CICS was reimplemented on Unix to be portable, it was implemented as a three-ball system built above Transarc's Encina toolkit. Now it is fair to say that all the popular TP monitors, such as IMS, CICS, ACMS, Pathway, Encina, Topend, and Tuxedo, are three-ball systems.

TP Monitors and ACID Transactions

Originally, *TP monitor* meant teleprocessing monitor—a program that multiplexed many terminals (clients) to a single central server. Over time, TP monitors took on more than just multiplexing and routing functions, and TP came to mean transaction processing. What's a transaction? A transaction is a set of actions that obeys the four so-called ACID properties—atomic, consistent, isolated, and durable (see the text box "The Four ACID Properties").

The ISO/TP standard (ISO0026) defines how to make transactions atomic, and the X/Open Distributed Transaction Processing standard defines a system structure and API that lets servers participate in atomic transactions. The transaction tracking system in NetWare, the resource manager interfaces in TP systems, and the transaction mechanisms of many SQL products can be used to build ACID applications.

TP monitors go well beyond a database system's narrow view of ACID applica-

tions. A TP monitor treats each subsystem (i.e., database manager, queue manager, and message transport) as an ACID resource manager. The TP monitor coordinates transactions among them. For example, a TP system will assure that when a database gets updated, an output message is delivered and an entry is made in the work queue, either all these actions will occur (exactly once) or none will.

Beyond ACID, TP monitors configure and manage client/server interactions. They help applications designers build and test their code. They help systems administrators install, configure, and tune the system and help the operator with repetitive tasks. They also manage server pools. Finally, they connect clients to servers and provide efficient system services to applications.

Queued, Conversational, and Work-Flow Models

Most TP monitors have migrated from a two-ball to a three-ball model in which the client performs data capture and local data processing and then sends a request to a router. The router brokers the client request to one or more server processes. Each server in turn executes the request and responds. Typically, the server manages a file system, database, or BBS shared among several clients. This design has evolved in three major directions: queued requests, conversational transactions, and work flow.

Queued TP is convenient for applications where some clients produce data while others process or consume it. E-mail, job dispatching, EDI (Electronic Data Interchange), print spooling, and batch report generation are typical examples of queued TP. TP monitors include a subsystem that manages transactional queues. The router inserts a client's request into a queue for later processing by other applications. The TP monitor may manage a

pool of applications servers to process the queue. Conversely, the TP monitor may attach a queue to each client and inform the client when messages appear in its queue. Messaging applications are examples of queued transactions.

Simple transactions are one-message-in, one-message-out client/server interactions, much like a simple RPC (remote procedure call). Conversational transactions require the client and server to exchange several messages as a single ACID unit. These relationships are sometimes not a simple request-response but rather small requests answered by a sequence of responses (e.g., a large database selection) or a large request (e.g., sending a file to a server).

The router acts as an intermediary between the client and server for conversational transactions. Conversational transactions often invoke multiple servers and maintain client context between interactions. Menu and forms-processing systems are so common that TP systems have scripting tools to quickly define menus and forms and the flows among them. The current menu state is part of the client context. Applications designers can attach server invocations and procedural logic to each menu or form. In these cases, the TP monitor (router) manages the client context and controls the conversation with a work-flow language.

Work flow is the natural combination of conversational and queued transactions. In its simplest form, a work flow is a sequence of ACID transactions following a work-flow script. For example, the script for a person-to-person E-mail message is compose-deliver-receive. Typical scripts are quite complex. Work-flow systems capture and manage individual flows. A client may advance a particular work flow by performing a next step in the script. The systems designer defines work-flow scripts as part of the application design.

Three-Ball Model: Routers Have X+A Connections

Clients Routers Application server pools Files

Routers manage pools of application-server processes and broker client requests to servers. This design scales by adding more routers.

THE FOUR ACID PROPERTIES

Transactions provide a simple model of success or failure. A transaction either commits (i.e., all its actions happen), or it aborts (i.e., all its actions are undone). This all-or-nothing quality makes for a simple programming model.

The ACID properties describe the key features of transactions:

Atomic All or nothing, either all the actions happen or none do.

Consistent The transaction as a whole is a correct transformation of the database.

Isolated Each transaction runs as though there are no concurrent transactions.

Durable The effects of committed transactions survive failures.

Database and TP systems automatically provide these ACID properties. They use locks, logs, multiversions, two-phase-commit, on-line dumps, and other techniques to provide this simple failure model. All the programmer need do is write consistent programs and bracket them with BEGIN and COMMIT. If anything goes wrong the programmer can call ROLLBACK—much like the quit function in a text editor. This simplicity is especially important in client/server computing and distributed databases, where the transaction may have done work at many nodes. COMMIT makes all the changes at all the nodes durable; ROLLBACK undoes all the changes.

Administrative tools report and administer the current work-in-process.

Writing Applications in a TP System

TP systems vary enormously, but the general programming style is to define a set of services that the server will provide. Each service has a message interface. To implement a system, you define client programs that generate these messages and server programs that service them.

In the two-ball model, your server program runs inside the TP monitor as, for example, a Sybase Transact SQL, Oracle PL/SQL, or CICS program. In the three-ball model, a C, C++, or COBOL program runs in a standard process using the TP system library to get and send messages.

Service programs are invoked with a message either from a queue or directly from a client. The service program executes the application logic and responds to the client with a response message.

TP systems generally provide tools to automate construction of forms-oriented clients and of simple applications servers that can serve as templates for customization. Other tools come from independent tool vendors, such as Intersolv, Magna, Texas Instruments, and Dynasty.

System Configuration, Management, and Operation

TP monitors allow you to build enormous client/server systems. Before three-ball TP monitors appeared on Unix systems, it was hard to scale up to (never mind manage)

300 clients. Today, most vendors are reporting TPC-A and TPC-C benchmarks demonstrating thousands of clients attached to a single server and tens of thousands attached to server clusters.

Now that you can build such systems, how do you manage them? TP monitors have evolved a broad suite of tools to configure and manage servers, as well as large populations of clients.

The TP monitor thinks of each database and application as a resource manager. These resource managers register with the TP monitor, which, in turn, informs them of system checkpoints and shutdowns and provides an overall operations interface to coordinate orderly system start-up and shutdown. Database systems, transactional queue managers, and remote TP systems all appear to be resource managers.

The TP monitor views each application as a collection of service programs running on one or more servers and provides tools to package these services relative to servers. Typically, short-running or high-priority services are packaged together, and batch or low-priority work is packaged in separate server processes.

After packaging the services, an administrator assigns security attributes to them. Aspects of security typically include role, workstation, and time, so, for example, clerks might be allowed to make payments from in-house workstations during business hours. The TP monitor provides a convenient way to define users and roles and to specify the security at-

tributes of each service. It authenticates clients and checks their authority on each request, rejecting those that violate the security policy.

Each service has a desired response time. Some services are long-running reports or minibatch transactions, and others are simple requests. TP monitors manage the size and priority of server pools so that each service meets its response-time goals. The administrator can control how many processes or threads are available to each service. Server pools can be spread over multiple nodes of a cluster, and this number can be dynamic, growing as loads do. When the number of requests for service exceeds the maximum size of the server pool, requests are queued until the next server becomes available.

Systems involving thousands of clients and hundreds of services have lots of moving parts. Change is constant, and TP monitors manage it on the fly. They can, for example, install a service by creating a server pool for it. Even more interesting, they can upgrade an existing service in place by installing the new version, creating new servers that use it, and gradually killing off old servers as they complete their tasks. Of course, the new version must use the same request-reply interface as the old one.

TP systems mask failures in a number of ways. At the most basic level, they use the ACID transaction mechanism to define the scope of failure. If a server fails, the TP monitor backs out and restarts the transaction that was in progress. If a node fails, it migrates server pools at that node to other nodes. When the failed node restarts, the TP system's transaction log governs restart and recovery of the node's resource managers.

Innovative TP Techniques

Modern database systems can maintain multiple replicas of a database. When one replica is updated, the updates are cross-posted to the other replicas. TP monitors can complement database replication in two ways: First, they can submit transactions to multiple sites so that update transactions are applied to each replica, thus avoiding the need to cross-post database updates.

More typical, TP systems use database replicas in a fallback scheme—leaving the data replication to the underlying database system. If a primary database site fails, the router sends the transactions to the fallback replica of the database. This hides server failures from clients—giving the illusion of instant fail-over. Because the

TP-Lite vs. TP-Heavy

According to database advocates, simple database transactions with stored procedures are all the transaction management anyone requires. That's the "TP-Lite" philosophy. In contrast, advocates of TP (transaction processing) monitors extend the notion of transactions beyond data resources. TP monitors track the execution of functions on a single server or several servers. That's the "TP-Heavy" approach.

TP-Lite is to TP-Heavy as a bicycle is to a Harley-Davidson motorcycle. TP-Lite can best be defined by what it lacks. TP-Lite functions don't execute under global transaction control. There's no global supervisor and only primitive process management. TP-Lite server functions incorporate only a single resource manager (i.e., the local database), and they don't support any form of ACID (atomic, consistent, isolated, and durable) nesting. TP-Lite may be adequate for a number of purposes, but the missing ingredients are rarely mentioned in the database marketing literature.

Scope of the Commit

A TP-Lite stored procedure is written in a database vendor's proprietary procedural language (e.g., PL/SQL and Transact SQL) and is stored in the database. A stored procedure is a transactional unit, but it can't participate with other transactional units in a global transaction. It can't call another transaction and have it execute within the same transaction boundary. So you have to write large transactions that put everything within the scope of the commit, which doesn't help the cause of modularization or reuse.

In contrast, TP-Heavy procedures are written using standard procedural languages. For TP-Heavy, global transactions are second nature.

Managing Heterogeneous Resources

A TP-Lite stored procedure can commit transaction resources that are only on the vendor's database or resource manager. It cannot synchronize or commit work that is on a foreign (local or remote) database or resource manager.

TP-Heavy procedures can easily handle ACID updates on multiple heterogeneous resource managers within the scope of a single transaction.

Process Management

A TP-Lite stored procedure gets invoked, executed under ACID protection (within a single-phase commit), and can then be cached in memory for future reuse. That's about it.

In contrast, TP-Heavy processes are pre-started and managed as server pools. As the load on a server pool grows, more processes start automatically.

Server pools support priorities and other class-of-service attributes. Server processes are protected from one another. If a server process dies, it can be restarted, or the transaction can be reassigned to another server process in that class. The entire environment runs under the constant supervision of the TP monitor. The server pool concept helps the TP monitor understand what class of service is required for a particular group of functions. It's an intelligently managed environment.

Client/Server Invocation

The TP-Lite stored procedure invocation is extremely nonstandard. Vendors provide their own proprietary RPC (remote procedure call) invocation mechanism. The RPCs are not defined using an IDL (Interface Definition Language), and they're not integrated with global directory, security, and authentication services. The communications links are not automatically restarted, and they are not under transaction protection. In addition, TP-Lite does not support MOM (message-oriented middleware) or conversational exchanges.

In contrast, the TP-Heavy environment is open to different communications styles. The RPC can use DCE (Distributed Computing Environment) as its base. Queued TP (a common option in TP monitors) is reemerging with the new name MOM. Most TP monitor vendors also support APPC/CPI-C (Advanced Program-to-Program Communications).

Performance

TP-Lite stored procedures are faster than networked static or dynamic SQL. But they don't perform as well as TP-Heavy managed procedures, especially under heavy loads.

The TP monitor offloads the database server by multiplexing client requests. It acts as a funnel on top of whatever funnel the database may have already put in place. So the TP monitor helps with both scalability and performance. Significant cost savings can be achieved because fewer database resources are needed to support a given workload. The Standish Group says: "Today, virtually all standardized transaction processing benchmark results (e.g., TPC-A) are executed with a TP monitor managing the application services in front of a database."

The Bicycle or the Harley

TP-Lite is useful in situations where you're dealing with a single vendor's database and fewer than 50 users. It's ideal in these entry-level environments, because you have to deal with only one server component: the database. The TP-Lite vendors also understand how to market their wares—an important advantage.

However, TP-Heavy technology is critical to the future of client/server computing. TP monitors allow us to mix components in all sorts of wild combinations; at the same time, they guarantee that everything meshes like clockwork.

router uses ACID transactions to cover both messages and database updates, each transaction will be processed exactly once. Without this router function, clients must explicitly switch servers when a primary database server fails and probably won't preserve exactly once semantics during the switchover.

RPC and request broker technologies are the bread and butter of TP systems. The emerging crop of commercial ORBs (object request brokers), whose performance can sometimes be measured not in

transactions per second but rather in seconds per transaction, would do well to embrace the three-ball model and the techniques pioneered by TP systems. Meanwhile, the venerable TP monitors, CICS, IMS, ACMS, Pathway, Tuxedo, Encina, and Topend, will continue to evolve as high-performance message routers and server-pool managers. ■

Jim Gray is a specialist in database and transaction-processing systems. He has worked for IBM, Tandem, and Digital on projects including System R, SQL/DS, DB2, IMS-Fast Path, Encompass, NonStopSQL, Pathway, TMF, Rdb, DBI, and ACMS. He is editor of The Benchmark Handbook for Database and Transaction Processing Systems *(Morgan Kaufmann, 1993) and coauthor of* Transaction Processing Concepts and Techniques *(Morgan Kaufmann, 1992). Jeri Edwards is director of transaction processing and client/server software development at Tandem Computers and one of the authors of* Essential Client/Server Survival Guide *(Van Nostrand Reinhold, 1994). You can reach them on the Internet or BIX at Gray@crl.com or Edwards_Jeri@Tandem.com, respectively.*

Conceptual design by Robert P. Humeniuk and Jerry D. Flynn of McDonnell Douglas Space Systems

Out-of-this-world graphics have *landed* on the *Intel* platform.

Personal workstations from Intergraph Computer Systems transport you to a world where high-end graphics software runs alongside your office automation tools — at a cost that won't send your budget into orbit. Until now, the processing power required for high-level CAD/CAM/CAE software forced you to work in two separate worlds: a PC for your office tasks and a workstation for intensive graphics design.

Now you can experience warp speed in both worlds. Personal workstations (TD-2 through TD-5) are equipped with single or dual Intel Pentium processors. In addition, they implement a workstation-like architecture that boosts Pentium power. So compute-intensive engineering operations — and your Microsoft Windows applications — run at lightning speed. And you can choose either Windows NT or Windows/DOS.

Light years beyond other systems, personal workstations are the first to implement the full thrust of OpenGL for graphics acceleration. So you can rocket through intensive 3D graphics operations such as rendering, modeling, and animation up to 100 times faster than conventional technology allows.

Why pay astronomical prices for workstations or push a PC beyond its limits? Choose the only Intel-based system made for the world of graphics — the personal workstation from Intergraph Computer Systems.

Call today: 800-763-0242 or 1-205-730-5441. Or reach us on the Internet at http://www.intergraph.com.

INTERGRAPH
COMPUTER SYSTEMS

Circle 244 on Inquiry Card (RESELLERS: 245).

MORE DISK SPACE.

Uninstall

FROM THE PEOPLE WHO BROUGHT YOU MORE MEMORY.

Have you ever tried to clean up your system? Then you know that removing all of the Windows™ apps, files and system components you don't want in order to make room for the things you *do* want is next to impossible.

Unless you have the right tool.

With Quarterdeck's new CleanSweep™ at your disposal, you can safely unclog your system and free up megabytes of precious disk space with the greatest of ease.

A DISK IS A TERRIBLE THING TO WASTE.

So before you blow a bunch of money on a new drive, grab your mouse and ask CleanSweep to analyze a program for you.

The long list of system files, DLLs, VxDs, Help files, .ini files, add-ons and other stuff that pops up on your screen will give you a pretty good idea of how much disk space you've been wasting. And since all those files are scattered all over

your system, they're *very* hard to find. Which is why CleanSweep comes equipped with exclusive Helper Technology that understands the bizarre installation habits of popular Windows programs.

TO DELETE. OR NOT TO DELETE.

CleanSweep knows more than just where to look for the maximum number of files. It also knows what to do when it finds them.

And what not to do.

Since uninstalling the wrong file–say, a shared DLL–can disable other programs you use daily, CleanSweep uses SuperLinks to track files that are needed by multiple applications.

It also lets you view any file, confirm every action, undo anything you've done and even

perform trial runs to see how much disk space you're about to gain.

Which is just the sort of ingenious functionality you'd expect from Quarterdeck, the folks who brought memory management to millions with the legendary QEMM.®

LEAVE NO FILE UNTURNED.

To really get the most out of CleanSweep, you'll want to use it to track down and delete duplicate, backup, archive, help and bitmap files.

Not to mention unneeded system components such as graphics drivers, fonts, wallpaper, system files and screen savers.

But before you can do any of this, you'll need to see your nearest reseller.

Or call us at **800-354-3222**.

Wherever you buy CleanSweep, if you're not thrilled, you can return it within 30 days for a full refund. No questions asked. Don't worry. CleanSweep uninstalls itself.

Quarterdeck®

Document Repositories

JONATHAN MACKENZIE

During the early 1980s, I spent six months prototyping a retail cash-processing application in Applesoft BASIC, and another six months with a compiler and a 6502 assembler honing the software to perfection. One day my boss called me into his office. "Look at this," he said, flicking through the menus and the all-too-familiar screens. "It's so simple." I flushed with pride; here was recognition at last. "Yes," I agreed. "It *is* simple." "So what the hell took you a year?," was the response.

Lotus Notes is a little bit like that: It has been evolving for nearly as long as Microsoft Windows has, and you can only guess how much was thrown away over the years. But what's remaining is deceptively simple.

Deceptively? Yes, because from an IT (information technology) perspective, you can suddenly build bigger and more distributed systems—with all the problems that process implies. It's also deceptive because such systems typically require conventional RDBMS (relational DBMS) capabilities that Notes awkwardly lacks. But for end users swamped with piles of documents, Notes makes it easy to build real client/server applications to help manage those documents.

Documents, Views, and Forms

A Lotus Notes database contains only documents—its data records. But these data records are miles ahead of the conventional kind (e.g., "Given *x* bytes per record, record N must be at offset N × *x*."). So, designer Iris Associates (which is now a part of Lotus Development) coined the term *document* for a dynamic data record of indeterminate size that can store any of seven different types of data: text, text lists, numbers, number ranges, times/dates, time/date ranges, and type composites.

A *type composite*, also known as a *rich-text field*, is a flexible, self-describing format that can store just about anything that you care to paste or embed into it. This includes word processor files, spreadsheets, bit maps, OLE objects, sounds, hypertext links (i.e., doclinks) to documents in the same or a different database, and even—if a C programmer using the Notes API so arranges—user-defined data types. To avoid confusion between a document in the Notes sense and a conventional document (e.g., a word processor or spreadsheet file), I will hereafter refer to

Notes data records simply as notes.

Views are windows for browsing through the database; they're sorted by pertinent data fields. Each column in a view either maps to a field in each note or shows the result of a Lotus 1-2-3–like formula. Using categories derived from the values in a database, a view can expand and contract in much the same manner as the outliner in a word processor does.

Forms are templates for entering and displaying a note's data. A form can contain static text, bit maps, data fields, and command buttons. The fields of a form map by name to the fields of a note, so a user can enter data with one form and view it using another.

From a user's perspective, you double-click on a Notes database icon to open a window containing a view. You then scroll through the view to the note that you want and double-click on it; another window opens to display the note's data in a form. It might contain, for example, an Ami Pro icon. Double-clicking on that launches the application and opens an embedded Ami Pro file for editing. When you quit Ami Pro, the file is stored safely back in the note.

Because notes are such flexible containers, a database can store all the elements of its own user interface within itself—icons, forms, views, macros, and so forth. Thus, each database becomes a completely self-contained application. *continued*

> **With Notes, even beginners can create and manage document repositories in true client/server fashion**

STEVE KELLER © 1995

Transmitting Documents by Replication

Replication means cloning a Lotus Notes application (or database—the terms are usually synonymous). There's no restriction on how many replicas can exist, and no single replica serves as the master. Each replica accumulates its own unique additions and deletions (the latter are represented by deletion stubs).

Replicas synchronize in two ways. *Server-based replication* is typically a scheduled affair; a task on one Notes server opens a dial-up or LAN connection to another. *Client-based replication* occurs intermittently at a user's discretion—perhaps from a laptop that dials in to a Notes server from a hotel room.

Either way, the involved parties first establish each other's identity. Next, they consult their replication histories and build a list of the data notes, design notes, and ACL (access-control list) notes that have changed since the pair last replicated. Subject to each application's access permissions from the other, they then start to exchange notes. The first note to be swapped is the ACL, which can then govern access permissions for the rest of the exchange.

After all notes have been successfully swapped, both replicas update their replication-history records. Since an application doesn't know which other applications it will replicate with in the future, or when, all deletion stubs are kept intact. They're finally removed after the deletion-stub-purge interval, which is typically 90 days.

What happens when, say, my coworker Fred and I use different replicas of the same database, which contains a note that has an embedded Ami Pro memo? I make some edits. Fred makes some different ones. Then we replicate with each other. (In a real-life situation, we'd likely both dial in and replicate with the office server, but the effect is the same.)

Sooner or later, however, a fight will inevitably develop. In fully connected client/server scenarios, user-interface software arbitrates overlapping edit requests. But Notes, only intermittently connected, handles conflicts the only way it can. Both notes end up in both databases, and the one that loses the editing battle is marked "Edit or replication conflict."

From then on, it's up to the humans to resolve things.

The Notes Document Repository

Some say that in any organization with more than 500 employees, there are no medium-size Notes databases. (And with its current average of 230 users per site, Notes isn't selling to small enterprises—yet.) A typical large organization has either one big application, with professional developers, thousands of users, and a budget to match, or a little database used by a team of about eight people, built by an IT-aware member of that team.

Typically, a large organization will start out with one big Notes application, often a personnel-oriented one with a hybrid client/server architecture. Such an application enables smaller workgroups, typically driven by the need to manage documents, to build and use their own Notes applets.

All organizations, large and small alike, have a standard document repository: the file server, a labyrinth of directories with thousands of files named in the DOS 8.3 format. Users can more effectively classify, navigate, and search for those files if they are placed under Notes' control. To create a simple Notes repository from the standard Document Library template, you only need about the same level of skill required to produce a useful spreadsheet.

Link, Attach, or Embed?

Word processor files can be left in native DOS format and linked to a note. Alternatively, there are three ways to move the contents of files directly into Notes—linking, attaching, and embedding.

Based on Microsoft's DDE protocol, *linking* is a kind of dynamic copy-and-paste operation that's just too clunky to be of any real value within an occasionally connected workgroup. I've seen grown men break down and cry trying to get the Paste/Special/Link menu option to work smoothly and consistently across a network. And if you're trying to work at home away from the LAN, forget it. That said, however, the two most useful protocols, DIP (Document Insert Protocol) and OLE, are both layered on top of DDE.

Any kind of disk file can be *attached* to a note. Notes stores a copy of the file in the document database, representing it as an icon in a rich-text field. DIP works by looking in the NOTES.INI file for the DDE strings required to launch an application and to pass it a command to open a file. The application need only be capable of being a DDE server. "DIP will probably go away by Notes version 3," a Lotus spokesperson once commented, "by which time most applications will be supporting OLE."

But today, on the eve of Notes 4.0, DIP is still with us and should not be underestimated as an effective tool in the document-repository designer's kit. Why? It has the merit of simplicity. When you double-click on a file-attachment icon, Notes asks whether you want to detach the attachment (i.e., extract a copy as a separate file) or launch it. If you choose to launch it, Notes creates a temporary disk file and then passes its name to the host application.

This technique works well when most material is read-only or is created by a few and read by the majority. But updating a file in place is awkward. You have to save it by another name, switch to Notes, put the note in edit mode, delete the existing attachment, attach the new file, save the note, and remember to delete the saved file from your disk.

By contrast, a document in the form of an *embedded* OLE 1.0 object is a lot easier to update. Like the attachment, a word processor file resides physically in the database. And, as with an attachment, the host icon is displayed in a rich-text field. But double-click on the icon, and you

During the edit of an embedded file, double-clicking on the icon in the rich-text field of a note launches the host application under OLE 1.0 control.

We've Made Network Faxing Easy and Affordable

Delrina has put the benefits of network faxing within your reach. Meet Delrina WinFax PRO for Networks™ 4.0 – the easy, affordable way to share fax modems and telephone lines across your network.

The benefits of WinFax PRO for Networks are enormous. Save the expense of buying individual fax modems and telephone lines. Make your people more productive, by letting them fax directly from their own PC. And, WinFax PRO for Networks is so easy to use and implement that you can have your entire group up and running in no time.

Easy to Use.
Network faxing is now as easy as printing from your PC. Using the world's best-selling, most award-winning fax software, WinFax PRO™ 4.0, you can send faxes directly from any Windows application. Then, choose any fax modem on your LAN or let WinFax PRO for Networks select one for you. To send a fax to a group of people,

New Delrina WinFax PRO for Networks 4.0

WinFax PRO for Networks will even use up to four modems at once. And because everything takes place in the background, you can continue working while you fax.

Affordable.
WinFax PRO for Networks is the most affordable solution for groups of up to 50 people. You don't need a dedicated fax server. Start off with a few users sharing one fax modem and grow until everyone in your department is faxing. To expand, simply add more fax modems – up to four per server and 16 per workgroup, with as many workgroups per LAN as you need. To add more users, simply purchase affordable 10, 25 and 50 volume user packs.

Compatible.
WinFax PRO for Networks also integrates with your existing hardware and software. It supports over 600 popular fax modems. It works with all popular networks including Novell NetWare,™ Microsoft Windows for Workgroups® 3.11 and Artisoft LANtastic™. Plus, it routes faxes using all popular e-mail packages like Microsoft Mail™ (MAPI), Lotus cc:Mail™ (VIM) and all MHS-compatible systems. If you don't have e-mail, simply forward faxes directly to another WinFax PRO user over your network.

Call Today!
The benefits of network-based faxing are now within your reach. WinFax PRO for Networks gives you an easy and affordable way to put faxing where it belongs–on your network.

Call now for information or to upgrade from WinFax PRO or see your local dealer.

N. America:	1-800-598-8679	U.K.:	44(0)181-207-3163
Germany:	+49/89/64 19 22-0	France:	33.1.46.94.80.40

DELRINA

Circle 242 on Inquiry Card.

Notes and Relational Data

In a conventional RDBMS (relational DBMS), a contact management application invariably stores company data in one table and contact information in another. Because there's a many-to-one relationship between contacts and companies, many contact records for a company can share a single company record.

But Notes is a document database, not a relational database, and it just can't work this way. Instead, applications designers typically rely on a useful feature of Notes called *named-field inheritance*. For instance, an application might present a company form with a button called Create New Contact. Its effect is to compose a new document that inherits the company name, phone number, fax number, and other fields, leaving just the contact name to be entered.

Unfortunately, this company information is now stored in the database twice. But other than their common data, no connection links

the two notes, so when you edit one, the other remains unchanged. Worse, if you change the organization note and then create another contact, the new note has the new address, but the older one does not.

One can argue that duplicating data is the only sane approach in a highly distributed, occasionally connected world. Nevertheless, it's often necessary to layer relational capability on top of Notes. There are three ways of doing this: You can build a custom Notes user interface to enforce relational integrity, build a server- or client-based agent to police the data, or tightly link Notes to a relational database.

A custom user interface for Lotus Notes is typically a Windows client program that's written in C and uses the Notes API. Alternatives include Visual Basic (with the right VBX [Visual Basic custom control]), Power-Builder, and ViP, the Lotus visual programming interface for Notes. To add relational

capability, a custom user interface must notice when a user saves a note so it can search the database for related notes and update them as well.

A server-based agent, scheduled to run at timed intervals, can achieve the same results. Written in C to the Notes API, it runs on the server as an add-in task. Of course, you have to compile the program with a C compiler that's native to the server operating system, dedicate a server for development and testing, and (if you want to use the add-in on a different platform) port it. Moreover, if the add-in program blows up while running on a production machine, it usually takes the server down with it.

A good alternative to the server add-in program is a Notes API program running on a client PC. A simple periodic task launcher can run an agent (or a series of them), wait a few minutes, and then start from the top again. If the task launcher and the agents take their configuration parameters from a Notes database on a server, the whole show can be controlled remotely (except for reboots) and the workstation can be locked in a cupboard.

A tight link between a Lotus Notes database and a SQL server usually takes the form of a program that's linked with both servers' API libraries (e.g., Notes and the Sybase DBLIB) and run by a scheduler task (see the figure "Synchronizing Notes and SQL"). Each time it runs, it must look first at one side (e.g., the Notes database) and ensure that all modifications made since the last time the database was visited are used to update the SQL server. Then it picks up all modifications made to the SQL database since the last time and uses them to update the Notes database.

Synchronizing Notes and SQL

Notes clients (usually read-only)

SQL server clients (usually read/write)

Server

9600-bps or faster modem

Notes server → Connecting agent → SQL server

A program linked to two APIs—that of Notes and that of a SQL server—can run on a scheduled basis to synchronize a Notes database with a relational database. Alternatively, it can be activated by SQL triggers or the Notes equivalent, DBHOOKVEC.

land in a new host application window with the file loaded (see the screen on page 132). Edit away, close the host application window, and the edited file is written back into the note. Since this is an in-memory operation, there's no temporary file to complicate matters. Lotus Notes 3.x supports OLE 1.0 as a client, not as a server.

For Unix users, there's LEL (Link, Embed, and Launch-to-Edit). Lotus Development and Software Pundits jointly designed and produced LEL at the same time

that Notes was ported to Unix. LEL is the functional equivalent to OLE 1.0 and is specially designed to let users store and retrieve documents in a mixed Windows/Unix environment. When an application is available on both these platforms, objects that are created on one platform can be launched and run by the application on the other platform.

At the time of this writing, only Ami Pro 3.0 enjoyed this exalted position. Like Microsoft, with its OLE toolkit, Lotus and

Software Pundits provide an LEL toolkit to enable Unix developers to incorporate the technology, both client and server, in their applications.

S. G. Warburg's Stream Repository

File attachments are a central part of an international marketing system, called Stream Repository, developed by the corporate finance division of investment bank S. G. Warburg. Jeremy Stone, the director responsible for the project, remembers

IN THE PIPELINE: NOTES 4.0

JON UDELL

Notes 4.0's first public showing, at Lotusphere in January, brought 6000 developers and business partners to their feet cheering. It should debut later this year. Listed below are just some of V4's many improvements.

A single document interface. Open a V4 database, and you land in a single three-paned window that shows a folder hierarchy, a list of items in the current folder, and a preview of the current item (see the screen). It's not a new idea, but it's a big improvement for Notes. In V3, each open document got its own independent MDI (Multiple Document Interface) window. These cluttered the screen and obscured the controlling view.

The current trend in efficient information display favors a single window with subpanes and sliding splitters. Lotus wisely jumped on the bandwagon.

Folders. For years, users and designers had to jam all database-level information display into collapsible-outline views accessed through the Views menu. In V4, views become folders, always visible—and navigable—in the folder pane. Folders in the view subtree collect documents automatically, based on a selection formula. Another subtree contains truly Mac-like nesting folders that hold ad hoc collections of documents. Other top-level folders hold the design elements and agents used to build V4 applications.

Web browsing. V4 can work as a native Web browser by means of on-the-fly conversion from HTML (Hypertext Markup Language) to Notes. A gateway server acts like a cache for a whole Notes network relative to the Web, so pages fetched by anybody on the network are available locally to everyone else.

Moreover, the local Notes database of Web pages can be fulltext-indexed and searched. If a page isn't locally available, it's fetched from the Web on demand, just as with Mosaic and Netscape. Pages are rendered as normal Notes documents.

If any Notes doclink contains URL (uniform resource locator) syntax, Notes will fetch the indicated Web page. Because the gateway connects to the outside world via TCP/IP, but to the internal Notes network only through the Notes API, it also acts as a kind of firewall.

A cc:Mail user interface. By popular demand, Notes mail will now look and feel very much like cc:Mail.

Location configuration. With Notes V3, it was a challenge to manage multiple replicas of a database where each was appropriate for use in a different location: the office LAN, home, and a hotel room, for example. V4 keeps all replicas of the same database as a stack of icons that appear as a single icon on the Notes desktop.

V4 also stores connection rules for each location. When you select your location, the appropriate icon bubbles up to the top, and that replica becomes current. You can customize replication for each replica, so the laptop version might replicate headers only while the LAN version replicates all fields, including attachments. And you can control the order of replication, so when time is short, you can be sure that your mail replicates first.

Improved formatting. The one-line-per-title constraint of V3 goes away; V4 supports multiline rows. Bullets, revision marking, and better table formatting are other welcome enhancements.

Viewlinks. When double-clicked, a doclink displays the referenced document. Analogously, a viewlink displays an entire database view.

Collapsible sections. Large documents can now be divided into collapsible sections. This technique aids the development of work-flow applications, since the parts of a document that are shown or hidden can vary on a per-user basis under programmatic control.

Subforms. V4 extends the inheritance capability of V3 with subforms, so documents can share common content within or across Notes databases.

A nonscrolling button bar. In V3, a button that invoked programmed actions appeared on the surface of a document and could disappear when a user scrolled through the document. V4 allows the use of a fixed region for action buttons, so they're always visible and accessible. The button bar can react to, and reconfigure itself for, changing values in database fields.

Graphical navigators. The default controlling metaphor for a database view is a folder hierarchy, but an alternative one is an image (a map, for example) with hot spots that, when clicked, drill down to the next level of detail.

LotusScript. V4 continues to support the Notes macro language, but it can alternately use the far more powerful and convenient LotusScript, Lotus's standard embedded programming language. LotusScript uses Visual Basic syntax with object-oriented extensions, features an integrated debugger, and can call components and DLLs—notably those that package the Notes API. In design mode, all the methods and actions available to the LotusScript programmer are displayed in browsers.

Actions. Triggering on such events as opening a document or performing a search, actions can take three forms: simple, formula, and script. A simple action requires no programming at all. You might, for example, hook the "send newsletter summary" action to a periodic search agent so that every search automatically mails a list of found documents. More complex actions can be programmed as formulas (i.e., Notes macros) or in LotusScript. Third-party applications can expose their own actions to Notes by means of an FX (field exchange) extension called NotesFlow.

A three-pane, Smalltalk-style browser packs a wealth of information and contextual clues into a single, efficient display.

Server management. V4 won't compete with NotesView, the new Hewlett-Packard OpenView-based management console for Notes, but it will offer baseline management tools to help administrators monitor server usage, analyze mail routing, and trace document flow.

OLE 2.0. As OLE 2.0 clients that support in-place editing, V4 document repositories can more seamlessly integrate host applications. LotusScript can drive applications that export OLE automation hooks. As an OLE automation server, V4 will be controllable from the outside by OLE automation clients. Some of the V4 user interface will be available for scripted control, as well as much of the API, according to Ray Ozzie.

Agents. V4 public agents are like V3 periodic macros, but they're server-based and can be written in LotusScript. They automate things that otherwise might be handled by server add-in tasks. Personal agents, which can run on workstations or servers, support per-user customization.

Pass-through authentication. V4 allows access to multiple Notes servers on a single call.

Faster replication. Field-level inheritance, and a faster method of exchanging replica IDs, should speed replication dramatically.

Scalability. V4 more aggressively exploits SMP (symmetric multiprocessing) hardware than V3 did and should also scale more effectively. It can run multiple replicator tasks concurrently on an SMP system, which should relieve bottlenecks in densely connected replication topologies.

Virtualization. A single V4 server can act like many independent virtual servers, each with its own set of users and databases. Why is this useful? With the advent of public Notes-based services such as AT&T Network Notes, it's handy to be able to divvy up server resources in this manner. A small company that can't afford to rent one of the boxes in AT&T's server farm will now be able to rent part of a server.

Jon Udell is a BYTE senior technical editor at large. He can be reached on the Internet or BIX at judell@bix.com.

how it got started. "Our worldwide marketing effort, much of it [involving] pitching to governments for the job of privatizing huge public-sector companies, relied heavily on hand-crafted, color-printed books presented to the people who made the decisions," he recalls.

"Yet all those books had to be produced by one or two centralized desktop publishing departments. The business is getting more marketing-intensive, and also turnaround times were getting shorter, creating a potential bottleneck. We decided that rather than increase the size of the central units to cope with increased volumes, we should take the opportunity to integrate the process of document production with all the other things that have to be done by corporate finance executives preparing for a transaction," he continues.

"We had to take the PC-based office software as given—Word and Excel were the group standards. Attaching key Word or Excel files to Notes forms enables us to control content and style and avoid [the need for] people reinventing wheels every day. The main components of the document repository are standard, frequently used items that can be freely mixed and matched in different documents but individually need to be controlled and kept up to date. The Notes database enables us to do this and also leverage the power of the tools we started off with," he concludes.

The principle behind Stream Repository is that a complete, street-ready presentation document can be assembled in Word with the aid of templates and macros. Graphs and diagrams are produced in Excel. A custom Chart Wizard, written in an Excel macro, allows the user to create graphs and place them on the clipboard to be pasted into the Word document. Lotus Notes databases are used at the beginning as a repository of frequently used, ready-formatted Excel charts, logos, and Word text paragraphs, and again at the end to store the completed presentation documents.

The simplicity of file attachment has worked well for the Stream Repository system. It has reduced potential problems in a combined Mac and PC environment as well as reducing support costs as the new technology has been rolled out worldwide.

Because the file attachments can't be indexed and searched by Notes' Verity Topic engine, the note that carries the attachment icon also has over a dozen different simple-text fields (e.g., Country,

Sector, and Type) that help users find what they're looking for as they navigate views of the database. In many cases, the note also carries keywords and even a short abstract to support fulltext search.

Coopers & Lybrand's CLASS

CLASS (Coopers & Lybrand Audit Support System), developed by the accounting firm Coopers & Lybrand, uses embedded objects. This system is a powerful implementation of Lotus Notes as a document repository.

As Rod Parry, the U.K. international partner for CLASS development and rollout, explains, "As you would expect, within C&L, the audit process is well understood, and it is, by definition, a workgroup activity. On a complex international audit, team communication is paramount, with perhaps 30 auditors in the U.S. and several groups of 10 in different locations around the globe.

"All of the team must have all of the audit file, all of the time. In addition, the system must enshrine our principle of TEQ [Total Engagement Quality] for each and every audit."

The CLASS system is already being used by 4000 people in 18 countries, with a planned 15,000 users in 25 countries by the end of 1995. It's split into three modules: Strategy, Audit File, and Audit Financials.

For each audit, the Strategy module, which is written in C, is used first. After an in-depth parameter-setting session, the Strategy module builds and populates a new Lotus Notes application, the Audit File, which is used as an audit plan, control-documentation store, on-line issue clearinghouse, and document repository.

Ami Pro and Lotus 1-2-3 are the word processor and spreadsheet used for the CLASS system. The templates for Ami Pro files and 1-2-3 spreadsheets are kept in a separate Lotus Notes database, called the Template Manager. This makes replication around the globe easier and ensures that all template files are up to date and fully meet legal requirements. Within the Template Manager, Ami Pro documents are embedded in a rich-text field, and a button on the form lets users copy the object to the clipboard.

When users switch back to the Audit File module with the insertion point in the required field, the Edit/Paste/Special/Embed menu option embeds the Ami Pro file in the note. As the audit proceeds, a complete database is built up, containing not only electronic copies of all letters, memos, and spreadsheets, but also all the

Special Report — Client/Server Computing

audit issues that are signed off as dealt with and all tasks that are signed off as done. Consistent quality is ensured for all audits, from one-person jobs to the largest international engagements.

Fulltext Searches

In practice, fulltext searches within a Notes document repository are rare. There are two reasons for this. First, attachments and embedded objects cannot be opened by the search if they are represented on a note by icons, although rich text pasted into a field is searchable. (This restriction is lifted in Notes 4.0, which can index attached and embedded content.) So, unless the host note carries a pertinent text abstract in a plain-text field alongside the object icon, it's invisible to the indexer and the search engine.

Second, a note usually carries other text fields (a letter, for instance, would have To, From, Title, and Date fields). Thus, it requires less thought, though a little more time, to find "that letter I sent to Jones last week about office equipment" in a view than it does to master even the simplest arcana of the fulltext search bar.

FX

FX, short for field exchange, is a valuable feature for document repositories. It allows the bidirectional exchange of data between a document, which is an OLE 1.0 object, and text fields in the note to which it belongs. Supported applications include Ami Pro 3.01, Freelance 2.01, Improv 2.1, Lotus 1-2-3 release 4.01, and, most recently, Microsoft Word 6.0c and Excel 5.0c (later versions of all these programs are also supported).

FX requires an OLE-embedded object in a note, SHARE.EXE (or VSHARE.386) to be loaded, and one or more hooks from the note to the object. In Ami Pro, for example, these hooks take the form of special bookmarks in the file (e.g., @TestField) and identical field names in the note (e.g., TestField). Note the leading @ for the bookmark name in Ami Pro.

When the object is launched, all fields that have matching names are read from the note and are used to update Ami Pro immediately before it displays the document. In a similar manner, when the host application window is closed, all the bookmark names are read and used to update the note.

With FX, for instance, the names of the fields in the note and their corresponding bookmarks in Ami Pro might be something like ToName, Organization, Address, DateOnLetter, Subject, and Title. If these bookmarks are edited while Ami Pro is running, the edits automatically update the fields in the note when the document is closed.

As mentioned earlier, a good Notes document repository is based on information stored in the note alongside the object. FX makes this job much easier. As you might expect, the technology used for FX is based on OLE; the object application is the server, and Lotus Notes is the client.

Further Refinements

It's possible to build automatic object activation into the design of a note so that when you open the note, the contained object launches. You can also control which actions—for instance, creating, reading, or editing of the note—invoke this behavior. Further, the note itself can be displayed or hidden.

When combined with FX, this produces the purest form of document repository. You open the Notes application, navigate a view to find the title that you want, and double-click on it. Your word processing application is then launched with the document in it. If you edit the title and close the document, the view will reflect the new title.

Because Notes lacks a distributed locking mechanism, it isn't a document manager in the purest sense. However, it provides a useful versioning mechanism. If you edit an embedded object within a note (or just edit a note), Notes saves a new note instead of overwriting the existing one, building a link between the two. This link ensures that, even if users recategorize on the fly, all the versions of a note are always clustered together.

Lotus Notes can support enterprise-document repositories, but it's also useful at the workgroup level. Users can store documents where they can easily find them again, use long names, categorize on any combination of fields, and painlessly replicate across continents, as well as to and from laptops.

This combination of features is certainly compelling. Lotus understands that Notes doesn't have to be the most powerful document manager or have the ability to satisfy all the criteria of, say, a professional librarian—it just has to be easy to use. And that it is. ∎

Jonathan Mackenzie, M.B.C.S., is a freelance consultant specializing in large, distributed systems. He can be reached on CompuServe at 100016,2107 or on the Internet or BIX at editors@bix.com.

Dimensions of Data

EDMUND X. DEJESUS

STEVE KELLER © 1995

Beth Ceja tracks statistical safety and accident information for over 17,000 employees at Southern California Edison. For two years, her tool of choice has been the multidimensional analysis package CrossTarget from Dimensional Insight (Burlington, MA). Reports come in from over 20 locations to a mainframe running DB2. Then they stage to the server-based multidimensional store that she analyzes using CrossTarget. Because such tools can rapidly "slice and dice" complex data in ways that are difficult or impossible for conventional SQL-oriented tools, they're an increasingly popular way to analyze raw numbers.

Data sources can include the following: mainframe databases, client/server relational databases, spreadsheets, data warehouses, text reports, flat files, and proprietary systems. Multidimensional tools can access file-based extracts, act as clients to SQL servers, or talk to specialized multidimensional servers. These OLAP (on-line analytical processing) servers and their clients are an important new flavor of client/server computing.

Flatland

Flat files, typical of legacy systems, are the most common data source. These files can be staged to a multidimensional structure on the mainframe itself, to a SQL server, or to a multidimensional server. Whatever the destination, a metadata layer (i.e., a data dictionary) is usually involved. Metadata describes the source, format, and layout of the extracted data.

Multidimensional products vary in their support for defining metadata. Some offer menu-driven or drag-and-drop interfaces to simplify the process (see the screen on page 140); others require considerable knowledge of source structures and transfer processes. Once in place, though, metadata remains fairly static.

Many vendors emphasize their products' ability to handle flat-file data, which ensures access to legacy data on mainframes. The flat-file format may also be the only option when transferring data between source and destination components that share no other compatible data formats.

Accessing Relational Data

It may not be necessary (or desirable) to replicate relational data to a natively multidimensional database. However, multidimensional tools that reach directly into SQL stores must create SQL statements to extract the data they analyze. As we saw with flat files, some products offer menu-driven or drag-and-drop interfaces to simplify the process. These interfaces enable non-SQL programmers to perform multidimensional analysis of SQL data. Because relational databases represent multidimensional reality the way a set of 2-D blueprints represents a house, defining metadata for access to relational data can be complex.

Relational database tables contain records (or rows). Each record consists of fields (or columns). In a normal relational database, one field in each record is the primary key, a unique identifier for each record. In contrast, the multidimensional data model is an *n*-dimensional array (sometimes called a Hypercube). Each dimension has an associated hierarchy of levels of consolidation of data. For instance, a time dimension might have a hierarchy with levels such as months, quarters, and fiscal years. Each location in a multidimensional array, corresponding to an intersection of all the dimensions, is a cell.

Variables (which are also known as measures or metrics) in a multidimensional array correspond to columns in a relational database table. Values within a table column correspond to values for that variable in a multidimensional array. Virtual variables (which are also known as off-page variables) in a multidimensional array are values that are not stored but calculated on the fly from stored variable values. An example of this might be variance, which is calculated from actual and predicted data. Multidimensional tools can manipulate and display virtual variables just

> **Multidimensional analysis in a client/server world**

as they do stored variable values.

Metadata maps the columns and rows of a relational database to the dimensions and cells of a multidimensional array. Metadata may also include rules for the consolidation of data at each level of each dimensional hierarchy. Variations on a time dimension, for example, might be

    days⇒weeks
    days⇒months⇒calendar years
    days⇒quarters⇒fiscal years

where data consolidates at each level of these hierarchies.

Multidimensional tools can access relational database data via SQL calls, possibly using multidimensional extensions to SQL that can, for example, treat columns as subscript-accessible arrays. The advantage of this approach is that multidimensional and relational tools can share a common relational data store.

However, there are also drawbacks. A SQL store is not a Hypercube. Instead, it's a set of related tables, and slice-and-dice analysis ("show me sales by region by month by salesperson") can be slow going, even if you've staged the data from the transactional server to a dedicated analysis server. Products that access relational data in this way include Brio Technology's DataPivot (used with its DataPrism), Cognos's PowerPlay, MicroStrategy's DSS Agent, and Trinzic's Forest & Trees.

Information Advantage has staked out a position firmly in the middle. The company's product called Decision Support Suite is an integrated set of applications to perform multidimensional analysis dynamically on a data warehouse and act as a server to a variety of front-end tools in an open three-tier client/server architecture.

Decision Support Suite is therefore a multidimensional analysis engine, accessing nonmultidimensional data and providing multidimensional information to client applications. This middleware approach does not dictate the form of the original data,

What the Users Say

HOW CISCO USES PABLO

Sales representatives for router manufacturer Cisco Systems have financial information at their fingertips, even when they're on the road. They use Andyne Computing's Pablo to access multidimensional data downloaded onto their Mac PowerBooks.

According to Ed Terpening, IS tools and technology manager at Cisco Systems, financial data (including transactions, sales, and revenue) resides in an Oracle relational database on a Hewlett-Packard Unix system. This data is loaded into a virtual data warehouse and then goes to the multidimensional structure on a Mac server. Pablo's multidimensional structure, called a Hypercube, contains data from the previous year (30 to 40 MB) that is updated weekly.

Terpening points out that, although about 30 to 40 people use Pablo now, he expects that number to eventually rise to about 100. They aren't just data analysts either. Pablo's ease of use allows comparatively nonexpert users, including regional managers and sales and marketing staff, to generate reports and comparisons as needed.

Because Cisco sales representatives compete against each other, each has only his or her own as well as aggregated data—no representative gets another one's private data. But having all that data resident, and having Pablo to slice and dice it, gives them a powerful advantage when they make sales calls.

HOW US WEST USES IRI'S EXPRESS

US West uses Express as both a source of data and a front end, according to Tom Keeney, director of financial systems development. Financial analysts in each business unit submit budgeting and other information directly, using

To prepare for multidimensional analysis, users graphically interconnect objects representing data sources, operations, and outputs. The example here is Planning Science's Gentium.

"Actually, we were quite lost without Axis print servers"

Printers in mixed environments are often a problem to both users and network managers. But it does not have to be like that.

Axis print servers, AXIS NPS 550 Risc and AXIS NPS 530/532, connect printers directly to your network, enabling all Ethernet benefits: the printer can be connected anywhere in the network, providing all users with access to high speed printing. Supervision and control over the printer is performed centrally.

In short: Network printing made easier.

NPS 550 RISC

UNIX

Novell

PC

Macintosh

IBM

Call now or fax this coupon for a free evaluation unit, tel +46 46 14 05 00, fax +46 46 14 05 10

Call from: France (16) 8532 6550, Germany 06404-91050, UK 01483-797 787.
Circle 575 on Inquiry Card.

and it allows access by a number of front-end tools.

Importing and Transforming Relational Data

Instead of reaching into SQL stores directly, some tools import relational data and transform it into a multidimensional array. It's a four-step process:

1. Establish the number of dimensions. Usually, it's the same as the number of keys in the relational database.
2. Traverse the key relations. Each logical group or collection represents a dimension in the multidimensional database. The multidimensional tool thus organizes the data into hierarchies.
3. Consolidate the relational database columns to appropriate levels. For example, data that is dimensioned by location may be consolidated into a hierarchy of

---

Structure of a Multidimensional Data Store

Array header

Array name,
Number of dimensions,
Number of cells,
comments

Dimension 1;
Number of positions
Pos1 name
Pos2 name
•
•
•
Fast offset calculation data

Dimension #n

Array Composition

Multiway arrays are comprised of two parts: a header and a data block. The header contains the basic data dictionary, or metadata, for the array, and the data block contains the actual values that populate the arrays. It is the combination of the header and data block that allows Multiway arrays to very quickly index and access large volumes of multidimensional data.

Array data

| True |
| 21.43 |
| Denver |
| 4 |
| $87.65 |
| 16 |
| Robinson |
| 3.14 |

Array Access and Manipulation

To access data, the header is used to determine which specific cell is requested. This allows the data to then be directly accessed in a single disk I/O. Because the header contains information about the size, shape, and attributes of each array, many dictionary changes (e.g., adding and deleting positions) can be done without ever changing the data, by accessing only the header. Moreover, new arrays can be created and manipulated through the header, reducing database maintenance and overhead.

Kenan's Multiway (part of the Acumate product) splits data into header and array blocks to streamline access. (Courtesy Kenan Systems Corp.)

geographical levels—district, state, and country.
4. Design the calculated data. Most tools let you design computed variables on the fly. *continued*

---

input screens running in Express. Monthly, a flat file is extracted from DB2 to the Express multidimensional database running on an IBM ES/9000 mainframe. The company is in the process of migrating this to a Hewlett-Packard Unix server. Users with Express running on PCs can access this multidimensional data and generate reports as they need to.

HOW RHONE-POULENC-RORER PHARMACEUTICALS USES IRI'S EXPRESS

Howard Mark, director of systems and programming for international Rhone-Poulenc-Rorer Pharmaceuticals, uses IRI's Express to track sales and marketing information worldwide.

According to Mark, every stable source of sales and marketing data (including prescription information, factory sales, promotions, sales forecasts, and sales calls) feeds into a mainframe collection point running Express. (Soon, however, Mark expects to replace the mainframe with a Pentium server.) Nightly batch runs populate the Express multidimensional database running on a LAN server. Some 50 users access the data, currently with Express EIS as a front end, though they may upgrade to Express for Windows.

HOW BLUE CROSS USES ESSBASE AND MICROSOFT EXCEL

Providers of health insurance have to track their own expenses as well as ours. For George Trudel, business and technology office consultant at Blue Cross/Blue Shield of Rhode Island, that means getting cost-accounting information off the mainframe and into Excel spreadsheets. His solution is to use Arbor Software's Essbase as middleware between the mainframe and the spreadsheets.

The expense data lives on an IBM mainframe, with updates flowing in daily. Each month, a flat file containing the past two years of this information (around 800 MB) is extracted to Essbase. Running on an IBM OS/2 server with 32 MB of RAM and a 2-GB hard drive, Essbase populates a database with dimensions, including time, cost center, line of business, and expense type. Excel macros can then pull data in from Essbase and fill the spreadsheets in about 5 minutes. (Trudel considered using Pilot Software's Lightship as a front end, but he found the spreadsheets adequate.)

Of special interest to Trudel was Essbase's support for multiple hierarchies. It can load data at the lowest level of each dimensional hierarchy. Users can then define rules for rolling up the data in different ways. Also crucial was Essbase's method of handling sparse data. The

Samsung
A new point of view.

SyncMaster 17GLs

It's about how you work...Attitude. Satisfaction. Pride. The 17-inch SyncMaster™ monitors reflect today's changing lifestyles. Hassle-free. Value-oriented.

Samsung took for granted you wanted bright, crisp clear images that don't flicker. You're also concerned about safety and energy conservation, so they're standard on the whole series. Life should be as simple as possible, so we added

SyncMaster	14GL	15GL	17GL	17GLs
CRT Size	14"	15"	17"	17"
Dot Pitch	.28	.28	.28	.26
Maximum Resolution	1024 x 768	1024 x 768	1280 x 1024	1600 x 1280
Maximum Refresh Rate at Maximum Res. (Hz)	76	76	60	60
MPR-II (Low Radiation)	•	•	•	•
Power Management	•	•	•	•
Flat Square Screen		•	•	•
On-Screen Programming			•	•
Real Color Control				•
Warranty	2 Years	2 Years	2 Years	2 Years

*1994 Samsung Electronics America Inc. The EPA Energy Star logo is not to be considered a product endorsement. Image on SyncMaster 17GLs courtesy of Mike Quon Design Office.

innovative new features like On-Screen Display and Real Color Control, and as you would expect, the GL Series comes with a two-year limited warranty.

Think about how you view your work. Samsung did. And came up with a whole new point of view. For additional information, call 1-800-933-4110.

TECHNOLOGY THAT WORKS FOR LIFE

SAMSUNG ELECTRONICS

The resulting multidimensional array may differ significantly from the source relational database. Extensive calculations, consolidations, comparisons, sorts, and selections can occur automatically during the transformation. If the relational database stores current operating data and historical data, the corresponding multidimensional array might create, calculate, and store future or projected data.

Because the data dictionary defines sources of data, dimensions and variables, dimensional attributes and their hierarchies, and formulas to calculate virtual variables or handle time series, it's a valuable by-product of transformation. This enhanced view of the raw data can greatly simplify conventional client/server applications development.

Going Native

Products that implement specialized multidimensional servers include Arbor Software's Essbase, Comshare's Commander OLAP Server, Holistic Systems' Holos, IRI Software's Express, Kenan Technologies' Acumate, Pilot Software's Lightship, and Planning Science's Gentium. These servers let users explore up to gigabytes of data, interactively teasing out meaning and insight.

Products that construct file-based multidimensional data stores include CrossTarget and Andyne Computing's Pablo. The vendors of such products claim greater portability of the data store (e.g., for use in mobile-computing situations), because the client/server architecture is not required to access the data.

Whether they're server- or file-based, multidimensional stores work with tons of data. How do they manage it? Sparse arrays are one key optimization. An array is sparse when few cells are populated or when many cells repeat the same value. Either way, the array can be compressed to cut down on storage and boost performance. Kenan's Acumate can be configured for different sparsity tolerances, and it splits the data into header and array blocks to streamline access (see the figure "Structure of a Multidimensional Data Store").

Arbor Software has patented a way to dynamically configure the dimensionality of the base unit of storage in Essbase to optimize its performance. This technique improves consolidation speed and memory utilization. It also cuts query response time and disk-storage requirements.

Predefining data location and format with metadata is also the key to rapid access. Some products, including Acumate,

database in question has over 4 trillion possible cells but is only 0.03 percent populated.

HOW AIRTOUCH USES ESSBASE AND EXCEL

Keeping track of your own phone bill is often a frustrating and mystifying experience. Imagine doing it for customers all over the globe, all in different currencies, and you have some idea of the challenges Don Sherwood faces. Sherwood, staff director of AirTouch International, a provider of international wireless communications services (including paging and cellular phones), has to consolidate information from 90 Excel spreadsheets to construct statements for customers and for internal use. After trying other solutions, he's settled on Arbor Software's Essbase.

Users have run their Excel spreadsheets for five years, and no one wanted to disrupt that part of the operation. It was important to preserve the software people were familiar with. Previous attempts to consolidate the information had included such diverse technologies as Next computers, combinations of Lotus 1-2-3 and FoxPro, and IFPS (Interactive Financial Planning System). However, the special rapport between Essbase and Excel clients makes this solution a natural fit.

"It's as simple as a drag and drop for us," said Sherwood. He selects the 90 Excel spreadsheet file icons and drops them on Essbase, running on a 486 OS/2 server. Essbase knows the structure of the Excel files. Using previously defined specifications, it can extract the required information (about 100 MB), perform calculations to transform currencies, and consolidate the information into a multidimensional database.

Sherwood's multidimensional database has six dimensions. Certain specific views are defined to automatically generate the needed reports.

Currently the system's sole user, Sherwood expects use to expand to five or more. The capability of the Essbase multidimensional database makes that expansion both attractive and feasible.

HOW SEARS USES ESSBASE AND LOTUS 1-2-3

Sears. Many stores. Many products. Many customers. Lots of data. How does it track actual sales versus projected sales? Irena Kleinaitis, retail reporting systems manager, consolidates mainframe and spreadsheet data with Arbor Software's Essbase multidimensional database server to generate a gross-profit matrix analysis.

Actual sales and inventory data reside on an IBM mainframe. Projections live in many Lotus 1-2-3 spreadsheets. Once a month, the mainframe data is passed as a flat file to Essbase, running on a 60-MHz Pentium PC with 256 MB of RAM and an 8-GB hard drive (though Kleinaitis points out that only 128 MB and 2 GB are needed). Data transfer, formerly a 39-hour marathon with a DB2-based product, now takes just 30 minutes. That improvement, and

2 WAYS TO PORT YOUR UNIX APPLICATIONS TO WINDOWS & NT:

HIRE MORE PROGRAMMERS

OR

TRY NuTCRACKER FREE FOR 30 DAYS

NuTCRACKER® is the fast and affordable way to port all your UNIX® applications to Windows™ 95 and Windows NT™ – without hiring an army of Windows programmers.

But don't take our word for it. We want you to use NuTCRACKER free for 30 days. Try it and you'll see how easy it is to leverage your investment in UNIX into fully interoperable native Windows applications.

NuTCRACKER is a complete set of UNIX development tools that looks and feels like UNIX to all of your programmers. It's the solution that's

already won enthusiastic converts at AT&T, Cabletron, Informix, and a host of U.S. government agencies.

With NuTCRACKER you can use your UNIX code and programming assets to develop new Windows products – right on Windows NT and Windows 95. It's also the best way to port UNIX applications across multi-CPU and multi-OS environments, without losing UNIX functionality.

NuTCRACKER also delivers a robust set of development tools. And if you have any questions, DataFocus offers everything from phone support to on-site Migration Services consulting.

So before you hire your own army

of Windows programmers, get all the details on the NuTCRACKER 30-day Free Trial Offer. Call or fax us today.

Tel: +1•703•631•6770
Fax: +1•703•818•1532

Cracking-Open Windows

NuTCRACKER from DataFocus Incorporated
12450 Fair Lakes Circle, Ste. 400, Fairfax, VA 22033
Phone (703) 631-6770, Fax (703) 818-1532
nutcracker@datafocus.com

Circle 264 on Inquiry Card.

provide direct cell-level access, which allows high-speed loading, updating, and retrieving of large quantities of array data. On suitable hardware, Kenan's ALI (Array Level Interface) can perform thousands of cell accesses per second, on both sparse and dense arrays. Such speed is necessary when populating a multidimensional database with gigabytes of information, even in an overnight process. These services are accessed through a C API that allows direct data exchange between application programs and the array.

Alternative Architectures

IRI Software's Express runs with the data engine and data dictionary residing on the server, while analytical screens and queries run on the client. A single database can be distributed among several locations.

Comshare's Commander OLAP Server includes a data refinery for data extraction, cleanup, and transport from legacy data sources. Part of the data refinery sits next to the source data on Windows NT, OS/2, VM, MVS, VMS, and Unix systems. Only the required data is extracted and moved, which cuts down on processing and improves performance.

Kenan's FocalPoint tool creates and manages high-level data links that afford "virtual" access. When you request data, FocalPoint checks to see if the data is in the multidimensional array. If this is not the case, it generates the code to access a relational database (using SQL) or a flat file (using file read commands). You don't have to know where the data is located or in what format.

FocalPoint loads data automatically at prearranged times, either statically (in batch mode, periodically) or dynamically (by user request to support on-the-fly drill-downs). The metadata defines the multidimensional data model, the locations of external data sources, the mappings between data sources and the array, and the distribution of data between sources and the array.

Multidimensional applications typically require access to multiple sources of data. Gauging the effect of sales promotions on sales volume, for example, may require tapping into separate advertising and sales databases.

Along with client/server support, multidimensional tools should handle local data storage and manipulation, so users can analyze data subsets off-line.

Open Access

Not everyone is in favor of multidimensional servers. Users rightly worry about

outstanding query performance, are welcome to the 40 users of the system, many of whom employ Lotus 1-2-3 as the front-end client to Essbase.

HOW NORTH MEMORIAL MEDICAL CENTER USES FOREST & TREES

Hospitals generate mountains of data on patients, so it's no surprise that administrators would turn to a multidimensional tool to help analyze it all. Ron Heim, vice president of information services and CIO of North Memorial Medical Center in Minneapolis, uses Trinzic's Forest & Trees to sift all that data for the information managers need. Heim's customers for this information are the president and the top two levels of management at the hospital. They require information daily about how many patients are using which of the hospital's locations for what services.

The patient information currently resides in a Sybase relational database running on Sun SparcStation 20s. (Laboratory results are on a separate IBM AS/400 and can be accessed via SQL calls that are transparent to the user.) Data is added to the patient database constantly, amounting to up to 10,000 patient records per day. Forest & Trees accesses the information every 30 minutes to provide snapshots of the system throughout the day. In addition, an overnight batch run sifts through over a million patient records so that Forest & Trees can generate daily reports to upper management.

The response time to the system is good, and Heim expects the current 20 users to swell to almost 60. The value of the system is that it provides timely views of how the hospital is serving the needs of its patients.

HOW PHILLIPS PETROLEUM USES SMARTSTREAM ANALYZER

Plastics, the party guest advised Dustin Hoffman in *The Graduate*. That advice has worked for the $1.8 billion plastics division of Phillips Petroleum. Keeping track of all the operational data involves plenty of computer power and the multidimensional capabilities of Dun & Bradstreet Software's SmartStream Analyzer.

According to Mark Evans, plastics IT (information technology) manager, the impetus behind the move to multidimensional analysis came from a senior executive who didn't like receiving filtered data. Using SmartStream Analyzer, managers can slice and dice raw data themselves.

Operational data includes product sales, inventory, manufacturing reports, and logistics (e.g., locations of rail cars). Data from a variety of mainframe legacy systems feeds a central repository implemented in DB2 on another mainframe. SmartStream Analyzer interrogates DB2 nightly and extracts data to a Sybase relational database on an IBM RS/6000. From there, the data is available for querying by over 30 top managers from their networked PCs.

The fact that managers can see operational data on many different levels

FIXED/REMOVABLE MASS STORAGE FOR ANY PLATFORM.

Data Express: Rugged Removability

Data Express, a family of durable removable carriers, houses a hard disk or DAT (Digital Audio Tape) device, adding up to 36GB plus the many benefits of storage removability to your PC or workstation. Data Express is available internally; mounting into your computer system drive bay, or externally; housed in steel enclosures and equipped with a fan and power supply. Data Express boasts an industry leading 25,000 insertions for long lasting removability. Constructed of steel and equipped with superior ventilation, Data Express provides the peripheral cooling needed when using large capacity devices.

Data Silo: Durable External Housing

Kingston's Data Silo is a family of stand alone external storage enclosures for half-height or full-height 5.25" or 3.5" SCSI peripherals. Data Silo is available in versions to house one, two, four, or nine SCSI devices simultaneously providing the utmost flexibility for storage expansion, disk array environments, and peripheral integration. Each Data Silo is equipped with its own power supply and fan and constructed of 100% steel, making Data Silo the most durable external storage enclosure available today.

Storage Versatility

Kingston's Data Silo four and nine bay units provide ideal drive stacking features for use in disk array environments. Data Silo also houses Kingston's Data Express products, which provide all the benefits of storage removability including data security and portability.

Industry-Leading, Five-Year Warranty

Every Data Express and Data Silo comes equipped with a comprehensive five year warranty and free technical support. Designed specifically for PC and workstation users, Data Express and Data Silo provide unsurpassed storage flexibility and quality.

Every Product 100% Tested

Kingston guarantees the highest quality available by testing every product prior to shipping.

Information At Your Fingertips

To get the facts on Data Express and Data Silo, call our convenient RAMFax fax on-demand service toll-free and request document number *8310*. For immediate assistance, contact Kingston's Storage Products group at:

RAMFax
CALL TOLL-FREE
(800)435-0056

(8 0 0) 4 3 5 - 0 6 7 0

Kingston
TECHNOLOGY CORPORATION
THE INSIDE NAME IN UPGRADES

Kingston Technology Corporation
17600 Newhope Street, Fountain Valley, CA 92708
(714) 438-1850 • Fax (714) 438-1847

Circle 246 on Inquiry Card (RESELLERS: 247).

proprietary access, the lack of tuning and maintenance tools, and the inability to update data incrementally. Many users will thus continue to rely on conventional relational systems for transaction-oriented work and supplement these with multidimensional systems for querying, reporting, and analysis. In fact, the Gartner Group predicts that natively multidimensional databases will fall behind products that optimize SQL-oriented databases for multidimensional analysis.

Vendors of multidimensional solutions are striving to make their products more open and accessible, embracing technologies such as ODBC (Open Database Connectivity), OLE, and DCE (Distributed Computing Environment). By publishing their own APIs, as Arbor Software has done with Essbase, vendors hope to attract third-party support and establish de facto standards in an otherwise standardless field.

Vendors increasingly stress the ability to access relational data on the back end and to display it in spreadsheets on the front end. Further, they try to serve up a variety of back-end data sources to a variety of front-end tools. Use of an existing client leverages the buyer's investment in software and the user's often considerable experience with the tool he or she knows best, so an architecture open to front-end clients is preferred.

The data dictionary is the crucial link between conventional data stores and multidimensional analysis tools. Most of the latter can already reach through and directly access original data sources. These capabilities will probably expand in the future. Some people predict the emergence of a standard MDQL (multidimensional query language).

Complementary Solutions

There's a trend toward partnership in the multidimensional camp. Comshare has licensed the Essbase engine from Arbor as a server component of Commander EIS. Microsoft Excel is strongly linked with Essbase. And Cognos's PowerPlay and Essbase are merging into an integrated OLAP solution. The establishment of the OLAP Council in January by charter members Arbor Software, Comshare, IRI Software, and Pilot Software is expected to produce a common API to simplify access (although some vendors complain they've been excluded).

Because OLAP and OLTP (on-line transaction processing) complement each other, the ultimate solution is to let both use the same data. Because OLTP must

can have unexpected consequences, says Evans. "A senior manager can call up lower-level management and say, 'Why has this rail car been sitting in Chicago for 120 days?'" Having used the system to great success in the plastics division, Phillips is considering it for the entire corporation.

HOW HERTZ USES COMMANDER OLAP SERVER

How can Hertz maintain a competitive edge against hundreds of competitors in thousands of locations worldwide? By consolidating and automatically monitoring competitive price data from multiple sources using multidimensional technology, according to Peter Ondi (director of global marketing planning information) and Bill Carroll (division vice president of marketing planning).

Data about competitors' prices flows in constantly from multiple sources, including travel agency databases and private providers. Hertz's mainframes provide its own internal data. On a daily basis, this multiple-source data is transformed into ASCII flat files and then filtered and merged with existing data into a multidimensional model.

A proprietary system of logic, rules, and criteria operates on the multidimensional model to automatically monitor prices and bring observations and trends of sometimes critical importance to the organization's attention. The model uses Arbor Software's Essbase running on a server. Comshare's Commander OLAP Server, running on Windows PCs, obtains multidimensional data from the model for analysis as a front end.

Carroll recalls that not too long ago, most of the data acquisition was done by phone and analysis was done by hand. Automated multidimensional analysis is both a dream come true and a business necessity. "To provide value to the customer, Hertz must be competitive on price," says Carroll. "This tool gives us a tremendous advantage with competitors in this market."

---

write and access data quickly, this argues for the data to be in a relational database. The problem then becomes how multidimensional tools can best access it. Ideally, there would be some way to rapidly transfer data from relational databases to multidimensional tools.

A big boost for multidimensional tools would be "data pumps" to move data from existing relational or nonrelational databases to native multidimensional databases automatically. Such data pumps already exist to move data to relational databases from other relational databases or from nonrelational databases. However, John Faig of the Meta Group sees a need for the ability of large organizations to automatically populate multidimensional databases in an off-line batch mode.

Relational databases and multidimensional arrays are not mutually exclusive, but complementary. Many organizations

will want to use both. There are several ways to configure the two systems to coexist peacefully. For example, information about the relational database can be included in the multidimensional server's data dictionary. This lets the multidimensional server access the relational database with dynamic SQL statements. Alternatively, tables in the relational database can be formatted to behave like extensions to the multidimensional server's data dictionary.

Someday soon, with luck, the details of OLTP/OLAP interoperability will be worked out. Users will then be able to get on with their real job: turning data into information. ∎

Edmund X. DeJesus is a BYTE senior editor. He has a Ph.D. in physics and has been a professional programmer for over 15 years. You can reach him on the Internet or BIX at edejesus@bix.com.

Client/Server with Distributed Objects

**ROBERT ORFALI AND
DAN HARKEY**

When Microsoft and CI Labs talk about OLE components and OpenDoc objects, they focus on the benefits of document-centric (as opposed to today's mostly application-centric) desktop computing. But these objects and components won't live only on the desktop. They're destined to become the currency of the intergalactic client/server era. How will OLE and OpenDoc talk to servers? How will object servers deliver the mission-critical support for transactions, concurrency, and scalability that today is the province of TP (transaction processing) monitors? The OMG (Object Management Group) has long anticipated these problems. Version 2.0 of the CORBA (Common Object Request Broker Architecture) aims to solve them by defining key object services, including transactions, concurrency, relationships, and externalization. Microsoft, with help from Digital Equipment, has a rival solution called COM (Common Object Model), and Novell's AppWare Distributed Bus may be yet another contender. All this enables us to say that we have seen the future, and it is distributed objects.

What's a distributed object? A C++ or Smalltalk object encapsulates code and data, and can be specialized by means of inheritance, but can't reach across compiled-language or address-space boundaries. In contrast, distributed objects are packaged as binary components accessible to remote clients by means of method invocations. Clients don't know which language or compiler built a server object, or where on the intergalactic network the object physically resides. They need to know only its name and the interface it publishes.

Ultimately we want supersmart client/server components that can not only interoperate but collaborate. For example, agents roaming the intergalactic network should be able to negotiate with other agents. In a homogenous programming environment, such as Telescript, that's easy. But the emerging world of software components is heterogeneous, and we need standards that set the rules of engagement among different types of components.

OMG's Object Management Architecture

The OMG envisions a common interconnection bus (see the figure "CORBA Standards for Component Evolution") that hosts client components, core services needed by all components (including naming, persistence, events, and transactions), and common facilities for component collaboration. These may be horizontal in the case of user-interface or system management facilities that interconnect components from different application domains, or vertical when specific to particular domains—say, car rental or hotel management. In these vertical domains, users will manipulate suites of business objects on the desktop, and tap into underlying client/server webs.

In the fall of 1990, OMG first published the *Object Management Architecture Guide*. It was revised in September 1992, and the details of the common facilities were added in January 1995. The four key elements of the OMG's architecture are:

- **The ORB (Object Request Broker)**: It's the object interconnection bus. Clients are insulated from the mechanisms used to communicate with, activate, or store server objects. CORBA 1.1, introduced in 1991, defined the IDL (interface definition language) and APIs that enable client/server object interaction within a specific implementation of an ORB. CORBA 2.0 specifies how ORBs from different vendors can interoperate.
- **Object services**: Packaged as components with IDL-specified interfaces, these services extend the capabilities of the ORB. OMG has adopted the following object services: naming, event notification, persistence, life-cycle management, transactions, concurrency control, relationships, and

> CORBA 2.0 mechanisms and services define how today's desktop objects and components can ride on tomorrow's intergalactic software bus

STEVE KELLER © 1995

externalization. Five new services are expected by mid-1995: query, licensing, properties, security, and time. These object services should be bundled with every ORB.

• **Common facilities**: These collections of IDL-defined components define the rules of engagement for application objects. They're categorized as horizontal and vertical. The horizontal ones address four disciplines: user interface, information management, systems management, and task management. The user-interface services, like OLE and OpenDoc, govern on-screen activities such as in-place editing. The information-management services resemble the OLE and OpenDoc mechanisms for compound document storage and data interchange. Systems management services define interfaces used to manage, instrument, install, configure, operate, and repair distributed objects. Task management services include things like work flow, long transactions, agents, scripting, and rules. In the realm of vertical facilities, IDL-defined interfaces will support suites of interacting objects specialized for health, retail, finance, and other domains.

• **Application objects**: These are components specific to end-user applications. To participate in ORB-mediated exchanges, they too must be defined using IDL. Application objects, of course, build on top of services provided by the ORB, common facilities, and object services.

The CORBA 2.0 ORB and object services should be incorporated into commercial ORBs starting in mid-1995. The application objects and common facilities,

CORBA Standards for Component Evolution

A common interconnection bus, the ORB (Object Request Broker) hosts systemwide object services, such as naming, persistence, and transactions; application-specific services, including user-interface technologies similar to OLE and OpenDoc; and application objects that wield all these other services.

however, are works in progress. When the common facilities are ready for prime time, perhaps by the end of 1996, CORBA will provide IDL interfaces for virtually every distributed service we know today. Note that it is not the OMG's goal to reinvent all of these services, but rather in many cases to provide IDL wrappers for existing standards.

The Structure of a CORBA 2.0 Object Request Broker

CORBA clients invoke services statically by way of compiled IDL (interface definition language) stubs or on the fly using of dynamic invocation APIs. CORBA servers route client requests to object implementations analogously by means of static skeletons or at run time using the new dynamic skeleton interface. Clients and servers share the utility APIs of the ORB (object request broker) interface.

New with CORBA 2.0

Street Machine

Up and away on the Information-Highway!

Mobile Assistant I™ means wireless communication, navigation and computer power in a small unit. Just six keys on the keyboard connect you to all the data networks you need all over the world. You can concentrate on your business: Prepare for your next customer call, process the last order or use programming at a remote computer site. The system takes you there - wherever you want to go or whatever you want to do - interactive, on a large size screen. As an option you can tune into a radio station or listen to your favorite CD.

Mobile Assistant I™ – a multifunctional device for mobile communications, traffic and information management. It brings together all the latest technology such as radio modems*, satellite navigation and computer power to give you the best service you can expect.

Mobile Assistant I™ is intended for professionals in areas such as: Freight Carriers • Field Service Organizations • Field Sales Organizations • Rental Car Agencies • Police- and Fire-Departments • Emergency Road Service • Hazardous Materials and Security Transports • Public Transportation Systems.

Now you probably want to find out more about this. Just give us a call!

SOLID® COMPUTER
...the open systems solution pioneer!

* MODACOM/ARDIS-MOBITEX/RAM Support

Workstation and SOLIDsummator (Software) — Mobile Assistant I™ — Portable Computer — PDA

Germany Headquarter: Bruckmannring 32 • D-85764 Oberschleißheim • Tel: + 49-89-31 57 19-0 • Fax: + 49-89-3 15 91 46

Circle 568 on Inquiry Card (RESELLERS: 569).

COM: The Other Object Bus

The COM (Common Object Model) is a distributed version of OLE 2.0's Component Object Model (also, and confusingly, called COM); it is Microsoft's alternative to CORBA. Digital Equipment, COM's cosponsor, is straddling both sides of the COM/CORBA fence with its ObjectBroker ORB. The COM specifications are preliminary; they will be finalized when Microsoft releases a beta version of its next generation of Windows NT, called Cairo.

Like OLE's COM, the new distributed COM specifies interfaces between component objects within a single application or between applications. Unlike the current COM, the new COM will provide local/remote transparency between components across networks, building on top of the DCE RPC (Distributed Computing Environment remote procedure call). Like CORBA, both COMs separate interfaces from implementations. Like CORBA, both COMs require that all interfaces be declared using an IDL (interface definition language). Microsoft's IDL, based on DCE, is not CORBA-compliant.

Unlike CORBA, COM does not support IDL-specified multiple inheritance. However, a COM component can support multiple interfaces. COM clients call QueryInterface at run time to discover which group of interfaces a component supports. Using a technique called *aggregation*, developers can package components within other components and can expose contained components' interfaces directly to callers of a containing component. Microsoft believes that aggregation and containment are all the reuse that's needed in distributed

environments. In the future, COM's QueryInterface may be coupled with a distributed authentication mechanism so that a component can express varying capabilities (interfaces) at run time according to clients' privileges.

Unlike CORBA, a COM interface is not a class in the object-oriented sense. COM interfaces do not have state and cannot be instantiated to create a unique object. A COM interface is simply a group of related functions. COM clients retrieve a pointer to access the functions in an interface—this pointer is not related to state information. COM may eventually use the OLE moniker concept to allow instantiation of named objects in the object-oriented sense: A client could use the pointer to reconnect to the same object instance with the same state (not just another interface pointer of the same class) at a later time. As described in the COM specification, monikers provide a combination of services, including naming, persistence, relationships, query, and object location.

COM, like CORBA, provides a DCE-based mechanism for generating globally unique IDs for its component classes and their interfaces.

COM, like CORBA, provides an event notification service through an interface called IAdviseSink. It enables components to register an interface that gets invoked when an IDataObject changes. OLE controls define a more generic event mechanism called *connectable objects*.

In this scheme, components specify (by means of IDL) the events they emit. Sink components, such as event handlers, can subscribe to these events. COM components support dynamic invocations by way of OLE automation, through an interface called IDispatch that is similar in concept to CORBA's dynamic skeleton interface. The current version of IDispatch (OLE 2.0) does not support networking.

We expect that the new distributed COM will eventually provide the same capabilities as CORBA, including alternatives to CORBA's persistence, transaction services, common facilities, interface repository, and relationships. The best we can hope for, short of Microsoft abandoning COM for CORBA or vice versa, is a single two-way gateway specification between CORBA and COM. This is exactly what OMG, with help from Microsoft, is trying to accomplish. Interoperability between OLE's nondistributed COM and CORBA is a snap; implementations are available from IBM, Iona, Candle, and Digital. Distributed COM is harder, because it's not yet fully specified. Still, it's clear that COM and CORBA will interoperate. But with dissimilar object models, components won't collaborate as effectively across the gulf as they can within each camp. Each model will have its own rules of engagement. Component providers will have to choose between OpenDoc/CORBA and OLE/COM on the client, and between CORBA and COM/Cairo on the server.

The ambitious goal of CORBA is to turn everything into nails, and give everyone a hammer. The nails are the IDLized services, and the hammer is the IDL interface to these services.

The ORB Client

A CORBA ORB looks more complex (see the figure "Structure of a CORBA 2.0 ORB") than it really is. The key is to understand that CORBA, like SQL, provides both static and dynamic interfaces. Let's first review how CORBA works on the client side.

Client IDL stubs provide the static interfaces to object services. These precompiled stubs, generated by the IDL compiler, define how clients invoke corresponding services on the servers. With dynamic invocation APIs, on the other hand, you can (at run time) discover a service that you want to invoke, obtain a definition of it, issue a parameterized call to it, and receive a reply from it.

Interface repository APIs allow you to obtain and modify the descriptions of all the registered component interfaces, the methods

they support, and the parameters they require. CORBA calls these descriptions *method signatures*. The interface repository stores, updates, and manages object interface definitions, and your programs use its APIs to access and update this information.

With CORBA 2.0, ORBs provide global identifiers called *repository IDs* that uniquely and globally identify a component and its interface across multiple ORBs and repositories. The repository IDs are system-generated unique strings that enforce naming consistency across repositories—no

Make CD-ROMs at your desk.

Multimedia publishing. Data backup. Software distribution. Electronic publishing. Records archiving. With the dataDisc series of CD-recorders you can put any information on a worldwide compatible CD-ROM and retrieve your information from any CD-ROM player.

These third generation CD Recorders will cut an industry standard 680 megabyte blank disc in less time with more reliability than comparable CD-ROM recorders. Plus, as the largest CD-R blank media retailer, we offer the best blank media at the best price.

Our systems come bundled with everything you need to begin making CDs right out of the box. With DiscMaker,™ our powerful CD mastering software, you can cut discs in **all** formats including CD-ROM (normal data), CD-DA (Digital Audio) and multisession format. Or add DiscMaker™ Multimedia mastering software and make CDs for XA (audio/video interleaving), CD-I, Photo CD, CD-TV and now Video-CD.

Fast, easy and affordable. The dataDisc CDR-2x comes bundled with DiscMaker™ mastering software, SCSI card, cable, 10 blank discs and lifetime, toll-free technical support for only $2,495.

Other configurations and faster CD-recorders are also available (CDR-4x and CDR-6x). Your data deserves the best - the best recorder, the best media, the best price. Call the number one CD-ROM company to get yours. Call dataDisc at 800-328-2347. Call today.

dataDisc™
We simply do it for less.℠

name collisions are allowed. They're generated by means of DCE (Distributed Computing Environment) UUIDs (universally unique IDs) or a user-supplied unique prefix attached to IDL-scoped names.

The ORB interface (used by clients and servers alike) supplies miscellaneous useful APIs. For example, there are APIs to convert between object references and strings. These calls are handy if you need to store or communicate object references.

With support for both static and dynamic client/server invocations and an interface repository, CORBA is more powerful and flexible than first-generation middleware, such as RPC (remote procedure call) (see "ORB vs. RPC" on page 160). Static invocations are easy to program, fast, and self-documenting. Dynamic invocations, while harder to program, offer maximum flexibility and are essential when applications must discover services at run time.

The ORB Server

Servers can't tell the difference between static and dynamic invocations. The same message semantics apply in both cases. The ORB locates an object adapter, transmits the parameters, and transfers control to the object implementation through the server IDL stub (also called a skeleton).

Static skeletons provide interfaces to each service exported by the server. These stubs, like the ones on the client, are created using

Is an Object Adapter a TP Monitor?

How does a CORBA server deal with millions of small objects? Who manages the object references? How are these references stored persistently? In theory, the answer to all these questions is the object adapter, CORBA's version of a TP monitor for objects. An object adapter registers the application's classes, instantiates new objects, gives them unique IDs, advertises their existence, invokes their methods when clients request it, and manages concurrent requests for their services (see the figure "The Structure of a Typical Object Adapter"). More sophisticated object adapters will also provide trans-

action management, load balancing, and fine-grained security.

An object adapter defines whether activating an object means creating a new process, creating a thread in an existing process, or reusing an existing thread or process. Servers may support a variety of object adapters to satisfy different types of requests. For example, an object database may want to implicitly register all its fine-grained objects without issuing individual calls to the object adapter. In such a case, it doesn't make sense for an object adapter to maintain per-object state, so the object database can provide a special-purpose

object adapter that interfaces with the ORB core and at the same time meets its own special requirements.

OMG prefers not to see a proliferation of object adapters. So CORBA specifies a BOA (basic object adapter) that can be used in many cases. CORBA requires that every ORB supply a BOA.

Today's object adapters aren't on par with TP monitors, but we expect this to change now that the object transaction and concurrency services have been adopted. TP monitor vendors may be the first to offer massively parallel object adapters that can orchestrate millions of server objects.

The object adapter:

❶ Registers server classes with the implementation repository. You can think of the implementation repository as a persistent store managed by the object adapter.

❷ Instantiates new objects at run time. The number of instances created is a function of demand for services. The adapter is responsible for balancing the supply of objects with the incoming client demands.

❸ Generates and manages object references. The object adapter assigns references (unique IDs) to the new objects it creates. It's responsible for mapping between implementation-specific and ORB-specific representations of object references.

❹ Broadcasts the presence of the object servers. The object adapter may broadcast the services it provides on the ORB, or it may respond to directory queries from the ORB core. Either way, it's job is to tell the outside world about the services it manages. Eventually, we expect to see tight integration with global directory services, such as X.500.

❺ Handles incoming client calls. The object adapter interacts with the top messaging layer

The Structure of a Typical Object Adapter

of the ORB communication stack, peels off the request, and hands it to the interface skeleton. The skeleton interprets the incoming parameters and presents them in a form that's acceptable to the object's method invocation.

❻ Routes the up-call to the appropriate method. The object adapter is implicitly involved in the invocation of the methods described in the skeleton stubs. It may be involved in activating the implementation and can authenticate incoming requests.

WE SAY IT'S EASY, AFFORDABLE NETWORKING. BUT DON'T TAKE OUR WORD FOR IT.

Whether your world is OS/2, Windows or DOS, LANtastic is the easy-to-use networking choice for small and medium businesses.

LANtastic makes networking easier than shaking hands, and provides you a flexible, functional and future proof solution with a price performance ratio that outperforms the rest.

LANtastic can expand rapidly, too. Start with as few as two users, add up to as many as a thousand.

To learn more about LANtastic, the network that makes your business a winner, ask any of our millions of satisfied users.

Of course, you can choose an easier solution: ask for the free Demo Disk and the free Guide to Easy & Affordable Networking by using the coupon, consult the nearest "Artisoft Premier" or "Artisoft Advantage" partner, or simply call 0753 554 999.

LANtastic®
BY ARTISOFT

Call now. 0753 554 999

©1995 ARTISOFT, INC. All rights reserved. Artisoft and LANtastic are registered trademarks and Artisoft Premier and Artisoft Advantage are service marks of Artisoft, Inc.

Circle 573 on Inquiry Card (RESELLERS: 574).

ASK OVER 2.4 MILLION USERS.

Free

YES, I want to know everything about the LANtastic network, and receive a free LANtastic 6.0 demonstration disk.

Tell me about: ☐ OS/2 ☐ Windows/DOS

Name	
Company	
Address	
City	Postcode
Tel.	Fax

Number of PCs to be connected:

☐ End User ☐ Reseller

BYTE 4/95

Return this coupon to: Artisoft. 223 Berwick Avenue, Slough, Berkshire, SL1 4QT, UK tel. 0753 554 999, fax 0753 551 325

CORBA's Object Transaction Service

With support for nested and inter-ORB transactions, it's an object foundation for dealing with the complex world of multistep consumer-to-business and business-to-business transactions

The CORBA 2.0 Object Transaction Service

Transactions are essential for building reliable distributed applications. So it should come as no surprise that transactions and distributed objects are getting married. OMG's (Object Management Group's) newly adopted OTS (object transaction service) defines IDL (interface definition language) interfaces that enable objects distributed across multiple ORBs (Object Request Brokers) to participate in atomic transactions, even in the face of catastrophic failure. It provides for interoperability between object transactions and procedural transactions that adhere to the X/Open DTP (distributed transaction processing) standard, and it optionally supports nested transactions. In short, it's an object foundation for dealing with complex, multistep consumer-to-business and business-to-business transactions.

An object involved in a transaction can assume one of three roles: transactional client, transactional server, or recoverable server (see the figure "The CORBA 2.0 Object Transaction Service").

● **A transactional client** issues a set of method invocations that are bracketed by begin/end transaction markers. The calls within the bracket may be for both transactional and nontransactional objects. You declare an object transactional in its IDL definition. The ORB intercepts the begin call and directs it to the transaction service, which establishes a transaction context associated with the client thread. The client then issues method invocations on remote objects. The ORB implicitly tags the transaction context and propagates it in all subsequent communications among the participants in the transaction. The

ORB also gets involved when the client issues a commit or rollback and notifies the transaction service. The client is oblivious to all this under-the-covers activity; it simply starts a transaction, issues its method invocations, and then commits or rolls back the transaction.

● **A transactional server** is a collection of one or more objects whose behavior is affected by the transaction but have no recoverable states or resources of their own. The ORB implicitly propagates the transaction's context whenever these objects call a recoverable resource. A transactional server does not participate in the completion of the transaction, but it can force the transaction to be rolled back.

● **A recoverable server** is a collection of one or more objects whose data (or state) is affected by committing or rolling back a transaction. Recoverable objects are transactional objects with ACID (atomic, consistent, isolated, and

durable) resources to protect (see the text box "The Four ACID Properties" on page 126"). They inform the transaction service that a recoverable resource has just joined the transaction whose context was propagated in the client call. Recoverable objects provide methods that are used by a transaction coordinator (i.e., the transaction service whose client originates the transaction) to orchestrate an ORB-mediated two-phase commit protocol.

OTS provides IDL-ized interfaces for the objects that make up the transaction service. So it is possible for clients and transactional objects to get more intimately involved in the details of the transaction via explicit method invocations. However, most transactions will depend on the ORB to transparently do all the work using its built-in facilities. Less intervention means better performance. OTS is a well-designed, low-overhead service that should perform at least as well as X/Open-compliant procedural transaction services.

an IDL compiler. A static skeleton provides hardwired support for the IDL-defined methods of a particular object class.

The dynamic skeleton interface, introduced in CORBA 2.0, provides a run-time binding mechanism for servers. The dynamic skeleton inspects the parameters of an incoming message to determine a target object and method. This technique is handy for building bridges between ORBs. It can also be used by interpreters and scripting languages that have to fabricate object implementations on the fly. Dynamic skeletons can receive either dynamic or static invocations from clients.

The object adapter sits atop the ORB's core communication services, accepting requests for service on behalf of the server's objects. It provides the run-time environment for instantiating server objects, passing requests to them, and assigning object IDs (object references) to them. The object adapter also registers the classes it supports

Generic ORB-to-ORB Bridging

Logical client-to-server operation request

Client — Full bridge — Server

Dynamic skeleton interface

Dynamic invocation interface

ORB A   ORB B

Logically, a client in communication with one ORB (Object Request Broker) just requests a service from a server associated with another ORB. Under the covers, ORB A uses the new DSI (dynamic skeleton interface) to route the client's request across an inter-ORB bridge, and then ORB B issues a dynamic invocation to the server.

and their run-time instances (i.e., objects) with the implementation repository. CORBA specifies that each ORB must support a standard adapter called the basic object adapter. Servers may, however, support more than one object adapter (see "Is an Object Adapter a TP Monitor?" on page 156).

The implementation repository lists the classes a server supports, the objects that are instantiated, and their IDs. It also serves as a common place to store trace information, audit trails, security, and other administrative data.

CORBA 2.0's Intergalactic ORB
CORBA 1.1 was only concerned with creating portable object applications; the implementation of the ORB core was left as an exercise for the vendors. The result was components that were portable but not interoperable. CORBA 2.0 tackled interoperability by specifying the IIOP (Internet inter-ORB protocol)—a scary name for TCP/IP with some CORBA-defined message exchanges that serve as a common backbone protocol. The messaging part of the backbone supports common data representations for all the OMG IDL types, interoperable object references, and common message formats and semantics optimized for ORB exchanges. Every ORB that

ORB vs. RPC

S o how are ORB (Object Request Broker) method invocations different from RPCs (remote procedure calls)? The mechanisms are similar, but there are important differences. With an RPC, you call a specific function (the data is separate). In contrast, with an ORB, you're calling a method within a specific object. Different object classes may respond to the same method invocation differently through the magic of polymorphism. Each object manages its own private instance data, so a method invocation affects data on a per-instance basis.

ORB method invocations have scalpel-like precision—the call gets to a specific object that controls specific data, and that implements the function in its own class-specific way. In contrast, RPC calls have no specificity. All the functions with the same name get implemented the same way. No differentiated service here. Of course, the ORB is often built on top of an RPC service, so you can end up paying a performance penalty for this extra level of service. (Note, though, that within a single address space, ORBs can dispatch methods almost as fast as C++ virtual function calls, so they can be useful for small objects as well as larger ones.) The ORB's overhead is worth every penny it costs if you're really exploiting the differentiation that distributed objects can offer. Otherwise you just bought yourself another layer of middleware, with all the costs and headaches that come with it.

calls itself CORBA-compliant must either speak IIOP natively or be able to bridge to it.

CORBA 2.0 specifies DCE as the first of many optional inter-ORB protocols, called ESIOPs (Environment-Specific Inter-ORB Protocols). At a semantic level, the DCE and Internet ORB implementations use the same messaging protocols. The IIOP is designed to enable out-of-the-box ORB interoperability over TCP/IP. The DCE ESIOP provides a

more robust environment for mission-critical ORBs requiring advanced features, such as Kerberos security, cell and global directories, distributed time, and authenticated RPC.

CORBA 2.0 also defines how to build ORB-to-ORB bridges (see the figure "Generic ORB-to-ORB Bridging Using CORBA 2.0's Dynamic Facilities"). When a client on one ORB calls a server on another ORB, the new DSI (dynamic skeleton interface) transmits the request to the target ORB, and then the dynamic invocation interface invokes the target object on that ORB. This dynamic technology is also well-suited for building gateways to non-CORBA object busses, such as COM (see "COM: The Other Object Bus" on page 154) and the Distributed AppWare Bus.

The CORBA 2.0 Object Services

The CORBA object services provide a unique way to build custom middleware. It's unlike anything today's conventional client/server systems offer. Component providers can use mix-in multiple inheritance to integrate their own code with the CORBA object services. Components gain persistence or transactional capability by inheriting these

things through IDL interfaces.

OMG currently defines standards for eight object services:

- **The life-cycle service** defines how to create, copy, move, and delete components.
- **The persistence service** provides a single interface for storing components in a variety of containers, including object databases, relational databases, and simple files.
- **The naming service** enables components to locate other components. The service allows objects to be bound to existing network directories or naming contexts, including ISO's X.500, OSF's (Open Software Foundation's) DCE, and Sun's NIS (Network Information Service).
- **The event service** allows components on the bus to dynamically register (or unregister) interest in specific events. The service defines a well-known object called an *event channel* that collects and distributes events among components.
- **The concurrency service** provides a lock manager that can obtain locks on behalf of either transactions or threads.
- **The transaction service** provides two-phase commit coordination among recoverable components using either flat or nested transactions (see "CORBA's Object Transaction Service" on page 158).
- **The relationship service** provides a way to create associations (or links) between components. It also provides mechanisms for traversing the links that group these components. The service can be used to enforce referential integrity constraints, and to track containment relationships and any other kind of linkages among components.
- **The externalization service** provides a standard way to stream data into and out of an object.

These services, and the five additional services OMG is currently working to define, empower components and provide a robust environment in which they can safely live and play.

Intergalactic client/server requires a common interconnection bus and an environment in which smart components can collaborate. CORBA 2.0 provides both and, with 440 member organizations backing the standard, we may finally get the universal plumbing we need to sustain the next client/server revolution. ■

Bob Orfali and Dan Harkey are (along with Jeri Edwards) the authors of Essential Client/Server Survival Guide *(VNR, 1994). They currently work on the application of distributed object technology at IBM. You can reach them on the Internet at harkey@vnet.ibm.com.*

Optiquest award winning quality raises the standard...not the price.

To some, a 'good buy' means having a low price tag. At Optiquest, we have a far higher standard. We build high-quality, technologically advanced monitors that are winning national and international quality awards AND we sell them at prices competitive to lesser quality monitors! Only Optiquest builds a complete family of fully EPA Energy Star compliant monitors each containing ergonomic and performance attributes that rival any competitive product. In fact, on October 12th, 1994 Optiquest was one of only 12 companies worldwide honored by the EPA as a 1994 Outstanding Energy Partner for 100% product line Energy Star compliance.

The new 17" 4500DC (15.7" diagonal viewing area) shown here features 1600 x 1280 non-interlaced resolution, advanced color management systems, fully digital controls, and programmable memory modes. Its street price will really surprise you. And as part of the Optiquest Complete Green

Solution, the 4500DC joins our whole line with valuable Opti-Green™ energy and environmental savings throughout.

To check out the award-winning quality, high technology, and very competitive pricing of the 14" to 20" Optiquest monitor family, just call 1(800) THE-OPTI. With all the good press we've been getting, you won't be alone.

Optiquest's Green Family Line

Monitor	CRT Size (Diag. In)	Dot Pitch	Maximum Non-Interlaced Resolution	Maximum Refresh	Controls Analog Digital	MAC Compatible	Color Matching	List
900S	14"	.28	1024 x 768(i)	87Hz	A			$ 265
1000S	14"	.28	1024 x 768	72Hz	A			$ 279
1500D	14"	.28	1280 x 1024	110Hz	D			$ 385
1562A	15"	.28	1024 x 768	76Hz	A			$ 389
2000DS	15"	.28	1024 x 768	100Hz	D			$ 419
2000DC	15"	.27	1280 x 1024	110Hz	D		Yes	$ 449
4000DC	17"	.27	1280 x 1024	100Hz	D		Yes	$ 765
4000TC	17"	.25	1280 x 1024	90Hz	D	Yes	Yes	$ 925
4500DC	17"	.28	1600 x 1280	100Hz	D	Yes	Yes	$ 859
2082DC	20"	.28	1600 x 1280	90Hz	D	Yes	Yes	$1,495

March 15, 1994
Optiquest 1500D

April, 1994
Optiquest 4000DC

June, 1994
Optiquest 4000DC

September, 1994
Optiquest 4000DC

October, 1994
Optiquest 4000DC

October, 1994
Optiquest 2000DC

October, 1994

OPTIQUEST®
Value Line Monitors by ViewSonic®
Circle 251 on Inquiry Card (RESELLERS: 252).

Break the Backup Barrier

Revolutionary Tape Backup at Four Times the Speed of Conventional Technology. RAPID-Tape Arrays.

10MB/sec! 36GB/hour! If your window of opportunity to backup your critical data is narrow, perform your backup with the fastest fault-tolerant solution available. RAPID-Tape arrays from ANDATACO incorporate industry standard 4mm, 8mm and Digital Linear Tape (DLT) tape drives and advanced Redundant Array of Parallel Independent Device (RAPID) technology, so data management is made simple. RAPID-Tape's resilient design allows your backups and restores to continue even when a tape drive needs service or a tape cartridge requires replacement. Experience the exhilaration of up to four times the capacity and data transfer rate of a single tape drive. RAPID-Tape, the fast and reliable backup solution that breaks all backup barriers.

- Compatible with all major operating systems and backup software utilities

- Appears to your system as a conventional tape drive. No device drivers required

- 560GB of unattended backup with RAPID Stacker systems

- 80GB capacity with an array of RAPID-Tape DLT drives

- Striping and parity to rapidly store and protect your data

- Simplify your disaster recovery and provide critical data security

- Mirroring allows for up to five copies simultaneously

- Hot-swappable user replaceable tape canisters and power supplies

- On-site installation, maintenance and dedicated 800 telephone support available

800-334-9191
Sales and Product Information

Shop Electronically
World-Wide Web URL:
http://www.andataco.com

ANDATACO
OUR REPUTATION IS ON LINE

The New Watcom SQL 4.0. You Would Think We Looked at Your Wish List.

Introducing the database server you've been waiting for: **Watcom SQL 4.0.**

With powerful features like stored procedures, triggers and cascading updates and deletes, Watcom SQL 4.0 can enforce your business rules and ensure data integrity.

In addition, dynamic multiple database support gives you and your users the unique ability to add and remove databases on the fly.

And, take a look at our self-tuning query optimizer; it learns as you use it. Each query performed against your database teaches the optimizer about your data, making it smarter... and faster. This dramatically reduces the costs associated with set-up and performance optimization of your DBMS by expert personnel. Also, the Watcom SQL product line includes both standalone and network servers built to allow you to move seamlessly between single and multi-user environments.

This is the database server with the functionality, performance and versatility you've been waiting for. Don't let the price fool you!

To check out the new features of Watcom SQL 4.0, call 1-800-395-3525 today.

Watcom
A Powersoft Company

One Box, Two Computers

With a PC inside it, Apple's DOS-Compatible Power Macintosh offers the best of both worlds

TOM THOMPSON

In the Spring of 1994, Apple Computer conducted an interesting market test. The company built an expansion board that was actually a fully functional PC computer. The board, code-named Houdini, plugged into the single slot of a Quadra 610 (see "Apple Provides PC on a Mac," BYTE, January 1994). Houdini had some design limitations: Its 25-MHz 486SX processor provided only modest performance, and it lacked support for NetWare operations and Sound Blaster I/O. Despite these flaws and the jaw-breaking name of DOS Compatibility Card for the Quadra 610, Houdini sold like mad, exhausting the inventory of 25,000 boards within several months.

Early this year Apple began shipping an updated Houdini that addressed the performance, networking, and sound I/O limitations that handicapped the original Houdini. Now we have Houdini's successor: DOS Compatibility Card for the Power Macintosh (DOS Card for short). As you might guess, this next-generation DOS Card fits in a Power Macintosh 601 PDS (Processor Direct Slot). You can obtain a DOS Card separately ($739) or packaged within a Power Mac 6100/66—the DOS Compatible Power Mac we review here.

Board Basics

The DOS Compatibility Card is a 7-inch expansion board with a 486DX2/66 processor. It supports 16-bit sound compatible with Creative Labs' SoundBlaster 16 board and 800- by 600-pixel SVGA graphics. DOS support comes via a Chips & Technology BIOS, and Apple bundles Microsoft's DOS 6.22 and Windows 3.11 with the board.

A single 72-pin SIMM socket on the DOS Card provides local memory expansion up to 32 MB. If this socket is empty at boot time, the Mac OS allocates the DOS Card a user-specified amount of RAM on the Power

The DOS Compatible Power Macintosh can run and display Mac and DOS sessions simultaneously. The DOS Compatibility Card with 486DX2/66 CPU (inset) provides much of the DOS capability, but Apple software makes it work well.

Mac's main logic board. In this mode, dedicated hardware transceivers perform the endian byte-swapping required to bridge the disparate memory organization the two processors use (little endian for the 486, and big endian for the PowerPC). A DMA channel that the Mac OS operates helps boost memory accesses in this shared mode.

So far, the DOS Card sounds like standard PC fare, but here the resemblance ends. The host Mac handles all serial I/O, disk I/O, keyboard, mouse, and network operations. For example, the Power Mac's floppy drive, which can read and write DOS and OS/2 floppies, fills in as the DOS Card's drive A. There is no parallel port I/O support, because the Mac lacks this interface.

Apple added code to the BIOS and wrote support programs to manage the peripheral I/O. A PC Setup Control Panel acts as the control center where you establish the data routes that the Mac uses to field the DOS Card's I/O requests. Here, for example, is where you might map the DOS Card's COM1 and COM2 serial I/O to the Mac's modem and printer ports. You occasionally have to watch out for Mac-specific limitations. Continuing with the PC serial port example, the maximum baud rate is 19.2 Kbps.

Across the Great Divide

The PC Setup Control Panel also sets the PC display resolution, the amount of memory the Mac will share (if the DOS Card

has none), sound I/O control, and drive C and D file mapping (more on this in a moment). You can also configure which hot-key to use for switching between the two systems, and for easier file exchange between DOS and Mac systems, you can set up a Mac folder that appears as a drive to the DOS Card. Support software lets you copy and paste certain graphics and text between the two environments.

You can set up file storage for the DOS Card in either of two ways: on the Mac drive as a "container" file or as a separate DOS-formatted SCSI drive. Container files are Mac files whose internal structure resembles a PC hard drive. To the PC Card, the container file appears as a DOS drive. From DOS, you can format volumes, erase files, and load all sorts of PC software. You can actually run DOS Installer programs that dutifully copy application software into this container file from floppies inserted in the Power Mac's floppy drive. Container files occupy the same amount of space as the equivalent DOS hard drive.

You can also mount a PC-formatted SCSI drive as long as you preformat it as a bootable device (FORMAT D: /s) on a DOS system. To make the drive accessible to the DOS Card, you select DOS partition from the pop-up menu in the PC Setup Control Panel's drive mapping section, and then select the drive's SCSI ID in the dialog box that appears. When you reboot both systems, the DOS Card readily accesses the drive partition. DOS access doesn't require special support files, drivers, or modifications to

the CONFIG.SYS file, which makes for real plug-and-play support for SCSI peripherals.

Doing Up DOS

We reviewed a DOS Compatible Power Mac—a Power Mac 6100/66 with 16 MB of RAM, a 500-MB internal SCSI drive, and an installed DOS Compatibility Card. The DOS Card came with 8 MB of local RAM. With the DOS Card and operating system preinstalled, set up came down to choosing the system's monitor arrangement: either a single-monitor configuration, where a hot key switches the screen between the two computers, or a two-monitor configuration, where you can simultaneously view both environments. Apple provides a scary-looking four-headed adapter cable with a 26-pin connector that plugs into the single connector on the DOS Card. The other connectors attach to the Mac video port, a Mac monitor, and to a PC joystick. This cable has a Mac DB-15 monitor connector, so you'll need an adapter to attach a VGA monitor to the system.

We preferred the two-monitor arrangement with both systems displayed. With the Mac busy on a CodeWarrior compilation, we could toggle to the DOS Card and check some article references via the network. You can use the shared mouse and keyboard with only one system at a time. However, when you toggle to the DOS Card, the Mac monitor dims, a nice visual cue as to which system owns the keyboard.

The Mac Extended Keyboard is a plus, because it sports the PC keyboard function keys, cursor keys, and numeric keypad, all of which operate as they should for a PC. The system software maps the Mac's Option key neatly to the PC's Alt key (labeled for both), which lets you operate and navigate both DOS and Windows application menus readily.

We connected a DOS-formatted hard drive to the Mac's SCSI port, and with a few mouse clicks and a reboot, it became drive D for the DOS Card. An external AppleCD 300 CD-ROM drive appeared both on the Mac's Desktop and as drive E to the DOS Card. The DOS Card was also able to connect to the BYTE network via the Mac

Ethernet port. Support for NetWare's IPX protocols is through an ODI (Open Data-Link Interface) driver supplied by Apple. This driver reroutes the IPX packets to the Mac's Ethernet driver.

You have to provide the remaining NetWare client software to complete the link. TCP support is also available through an ODI driver, but other network interfaces are not supported at this time. Once the AUTOEXEC.BAT file was configured properly, I was able to access all BYTE's NetWare file servers. The Mac OS was also able to access an AppleShare server and all networked laser printers simultaneously. As long as the DOS Card and Power Mac use different network protocols (e.g., IPX and AppleTalk, or TCP and IPX), both can share the Ethernet connection without trouble.

Compatibility and Performance

DOS and Windows compatibility was excellent. The DOS Card ran Doom—and even produced the sound effects—without problems. Under Windows, we could access a networked CD-ROM drive holding Ziff Communication's Computer Select. We could also use Folio Corp.'s Folio Bound Views application to search for items in the text of the last five years of BYTE, which was stored on a NetWare server. This setup allowed use of Mac applications, while providing easy access to resources that are available at BYTE only in a PC format.

In terms of performance, the host Mac has the kick of your typical 66-MHz Power Mac. As the BYTE Native Mode benchmark results show, the DOS Card's throughput nearly rivals that of a typical 66-MHz 468DX2-based PC. However, the comparison PC system had a 256 KB secondary RAM cache, which the DOS Card lacks. While the memory-sharing mechanism works flawlessly, it exacts a huge performance hit. On a DOS Card

---

NATIVE MODE BENCHMARK RESULTS FOR DOS

We tested the Power Mac DOS Card's low-level performance using BYTE's Native Mode test. Both integer and floating-point results were slightly slower than the comparison system, a 486DX2/66 PC. Without local RAM on the DOS Card, performance dropped markedly. The index score is based on a Dell 90-MHz Pentium system.

TEST	COMPARISON PC	DOS CARD WITH 8 MB LOCAL RAM	DOS CARD USING SHARED MAC RAM
Integer index	.415	.402	.244
Floating-point index	.265	.251	.143

---

sharing Mac memory, PC test results consistently plummeted by 40 percent. The Mac felt the hit less, perhaps because it physically owned the memory, but its performance degradation fluctuated between 13 to 40 percent, depending on the operation.

The Best of Both Worlds

As mentioned earlier, the DOS Card alone costs $739. A Power Mac 6100/66 with 16 MB of RAM, a 500-MB internal hard drive, and a DOS Card, but no monitor or keyboard, costs $2759. Neither price includes local memory for the DOS Card, but budget for it if you opt for just the board (you can move one of the Mac's RAM SIMMs to the DOS Card to obtain 8 MB of local RAM for each computer).

Overall, the DOS Compatible Power Mac makes a superb PC, especially with a Mac wrapped around it to handle the messy details of attaching peripherals. The Mac offers its convenience and graphics, and the DOS Card provides access to a wide variety of DOS and Windows programs. Although a PC in a Mac seems like a Jekyll-and-Hyde prospect, it actually delivers the best that each platform offers. ∎

Tom Thompson is a BYTE senior technical editor at large with a B.S.E.E. from Memphis State University. He is an Associate Apple Developer and author of Power Macintosh Programming Starter Kit *(Hayden Books, 1994). Contact him on AppleLink as T.THOMPSON, or on BIX or the Internet at tom_thompson@bix.com.*

About the Product

DOS Compatible Power Macintosh $2759
(16 MB RAM, 500-MB hard drive, DOS Card; price doesn't include monitor, keyboard, or DOS Card memory)

Apple Computer, Inc.
1 Infinite Loop
Cupertino, CA 95014
(800) 776-2333
(408) 996-1010
fax: (408) 974-6412
Circle 1002 on the Inquiry Card.

SoftWindows: The Next Generation

If you occasionally want to run DOS software on a Power Mac without buying a DOS Card, you might check out Insignia Solutions' SoftWindows (Insignia Solutions, 1300 Charleston Rd., Mountain View, CA 94043, (800) 848-7677). The company plans soon to release a major upgrade to its SoftWindows, a program that emulates the PC environment. SoftWindows 2.0 will emulate a 486 processor and thus offer 386 Enhanced Mode capabilities to those applications that require it. Performance gets a boost from the interpreter's dynamic compilation of frequently used x86 code segments. The new version also offers enhanced sound (through the Windows Sound System only, no Sound Blaster), SVGA support for DOS displays, and support for simultaneous use of more than one network protocol stack.

Visual C++ Goes Multiplatform

Microsoft's latest development system is a well-crafted set of tools for building Win32 applications—and it's not just for PCs anymore

STEVE APIKI

Visual C++ 2.0 is much more than an upgrade. With Microsoft's delivery of Alpha and Mips development environments, and especially with its shipment of the Cross Development Edition for Macintosh, Visual C++ has become a real multiplatform development tool. If you're a software developer concerned with targeting the Mac as well as various flavors of Windows and Windows NT, you now have a slew of new options.

On the other hand, if you are concerned solely with Win32, the choice is simple: Get Visual C++ 2.0. Enhancements to MFC (Microsoft Foundation Classes), support for 32-bit OLE controls, and a much-improved integrated environment make this an easy decision. Unfortunately, Win32-only is an unlikely scenario—for most developers, supporting the Mac, or 16-bit Windows, or even DOS are day-to-day realities. With add-ons, Visual C++ 2.0 can deliver on these platforms as well.

New and Improved
Version 2.0 succeeds both the 32-bit edition of Visual C++ 1.1 and Visual C++ 1.5 and adds significant improvements over the existing products. Chief among these are MFC 3.0 and a noticeably improved integrated environment.

MFC 3.0 extends the framework to incorporate virtually all Windows-related technologies. Version 3.0 adds classes that support ODBC (Open Database Connectivity) and OLE 2.0. (Although these were present in MFC 2.5, MFC 2.5 is a 16-bit class library, so Win32 developers had been left out in the cold.) MFC 3.0 further extends OLE capability with classes that model OLE controls (although not control containers), the 32-bit component replacements for VBXes (Visual Basic custom controls). New user-interface elements like tabbed dialog boxes, tool tips, and dockable windows, all of which Microsoft says will map directly to Windows 95 gadgets, round out the new classes. These and other MFC enhancements are fully supported through a new App-Wizard and a new ClassWizard.

The new integrated environment is a big improvement over those found in Visual

Integrated resource tools. External AppStudio is gone, replaced by integrated tools.

Full-featured integrated debugger meshes well with editor, letting you make edits while in the debugger.

New UI features like dockable windows and tool tips migrate to Visual C++ from Microsoft productivity applications.

Customizable toolbar with drag-and-drop access to buttons.

Syntax highlighting. In Version 2.0, syntax coloring is customizable.

Project management features allow control of build settings for multiple targets. This screen shows Macintosh targets, which are available only with the addition of the cross-development kit.

Incremental linking speeds links.

Tabbed output windows allow easy access to build results, find results, or output screen.

Find-in-files option lets you search for strings or regular expressions in multiple files without leaving the environment.

The New UI: Visual C++ 2.0 presents a significantly improved look, with a better integrated debugger and easier access to peripheral tools. This screen is from the Intel-hosted package; the core Mips and Alpha packages provide an identical set of options.

C++ 1.1 and 1.5. Most important, the debugger, resource editor, and other tools are more completely meshed within the IDE (Integrated Development Environment), and you aren't constantly launching additional tools. The IDE itself incorporates all the new user-interface features of MFC 3.0, so you get first-hand experience with these before you foist them on users. (As BYTE was going to press, Microsoft announced version 2.1, which includes some Windows 95 controls, an updated OLE developer's kit, and ODBC support for the Mips version.)

Moving Up to Win32
Generating a new application for Win32 is trivial with Visual C++: A trip through App-Wizard, where you select the Windows features you want to support, rewards you with a complete template from which you can bootstrap development. However, most projects don't start from scratch, so we evaluated version 2.0 by converting an existing 16-bit Windows application that was not

originally developed in MFC.

Moving from a 16-bit project to a 32-bit project is an all-or-nothing affair. Visual C++ 2.0 can't build 16-bit applications; 16-bit Windows support is relegated instead to a copy of Visual C++ 1.51 shipped on the same CD. Once you bring your 16-bit project into Visual C++ 2.0, it is forever after a 32-bit project, and you can no longer support Windows 3.1 without Win32s. This dichotomy between 16-bit and 32-bit is the single factor that might make a person hesitate to upgrade to Visual C++ 2.0.

To begin the conversion, we ran App-Wizard to generate a new MFC-based project. AppWizard and ClassWizard, as in previous versions, simplify source-code generation by never touching user code. This keeps your code safe, but it makes it impossible to "regenerate" an application should you want a new MFC feature later, or to modify the names of message handlers using ClassWizard. Therefore, this is a step you need to get right the first time.

continued

Converting an existing non-MFC project to MFC was, of course, the bulk of the work, but making the conversion is well worth the effort: MFC is at the right level to allow porting to other machines and other user interfaces like the Mac and Windows 95 when the time comes.

We confess to liking command-line tools and make utilities for flexibility and sheer variety. However, the level of integration in Visual C++ 2.0 is excellent, and new features like a Unix grep-like find utility make the IDE stand almost completely without external tools. Its weakest point is its editor, which is Mac-like in its command structure but can't emulate other editors.

With the editor, debugger, resource editor, and other tools so tightly coupled, you rarely have to leave the IDE. You can edit in the debugger or set debugger breakpoints in the editor. A double-click on a line from the results of a find-in-file command opens the file. AppStudio is gone, replaced by an integrated resource editor (which finally adds a text tool for the bit-map editor).

Visual C++ will not deliver the shortest build times you've seen, but the compiler thread can run in the background while you attend to other tasks. Version 2.0 also supports incremental linking, which really speeds up the link phase. One nice feature is the ability to build the browser database without completing a full link. While moving our project into MFC, we sometimes had conflicting definitions that kept the link from completing; the browser easily pegged them.

It took us a few weeks to get the core of our test application into MFC and up under Win32. The application also requires a DLL, which we also had to port to Win32. Once all that was up and running, we moved both projects over to Mips and Alpha versions of Visual C++ 2.0. Porting to both new platforms took less than an afternoon: The similarity among the environments is amazing. However, since the application was written to be portable, conditional compile statements had to be modified because they assumed that if _WINDOWS was defined, the program was being built for an x86 platform. There was also a slight difference in the va_list type between Alpha and other platforms that required making some modifications.

Targeting the Mac

Porting to Mips and Alpha NT was easy; porting to the Mac was not. Unlike other NT systems, which are both hosts and targets for Visual C++ 2.0, the Mac is supported only as a target through a cross compiler that runs on Intel-based NT systems.

The Visual C++ 2.0 Cross-Development Edition for Macintosh is a $1999 add-on to the standard Intel package. The package includes the cross compiler (which integrates into the regular environment), WPL (Windows Portability Libraries), a remote debugger, and utilities for transferring files and executing programs across an Ethernet connection. Once your 68K-based System 7.x Mac is hooked to your NT system running AppleTalk, you build programs on the NT system and ship them to the Mac, where the debugger can execute them remotely.

WPL is a set of static libraries that implement most of the Win32 API; you link these into your application and they mediate between Win32 calls and Mac Toolbox routines. The result is an application with a more-or-less native Macintosh look, although you must make some source and resource changes to complete the transformation.

The first step in porting our test project was to set up Mac targets within the development environment. This is simple to do, and you can set 680x0-specific code-generation features for these targets. Our test project included serial communications, which WPL doesn't support. We ended up leaving these features out of this stage. Finishing the port now required conditional Toolbox calls to replace the Win32 calls; the required Apple header files and libraries are included.

After taking out or replacing the Win32 API features WPL doesn't support, we built and linked the application. Visual C++ automatically downloads the program to the Mac, but you must remember to either include a size resource or set the memory partition size on the Mac before every run.

The test application requires a DLL under Windows. The cross-development tools support Apple's DLL implementation—ASLM (Apple Shared Library Manager)—allowing construction and use of shared libraries. This is a critical requirement, without which porting most real Windows applications would be impossible. However, converting a DLL to a shared library is difficult. Worse, the documentation on issues like what libraries are required is severely lacking, and the two technical support representatives we spoke to at Microsoft weren't versed in either Mac development or the cross-development kit.

Microsoft shows little confidence in the future of ASLM; WPL itself is a static library. The company recommends that developers also consider linking Windows DLLs as static libraries for the Mac wherever possible. The PowerPC version of the cross-development kit, due out before the end of the year, will support code fragments instead of shared libraries.

Building Mac applications under Windows NT limits tools on both sides of the connection. For example, because it is remote, the debugger requires that you set all breakpoints locally before beginning the debugging session. Also, some menus and dialog boxes from Windows are implemented as custom resources on the Mac, and you can't browse through them using programs like Apple's ResEdit or Mathemaesthetics' Resorcerer.

Living with Limitations

Despite having a few limitations, the cross-development toolkit is an excellent option if Windows or Win32 is your primary target. We would choose it over any existing cross-platform library for two reasons. First, if you already know the Windows API, you can generate a reasonable Mac application without learning an entirely new API. Second, it's a single development environment; you can build, run, and debug for both targets with just a click on each new target. In short, the Visual C++ 2.0 cross-development toolkit greatly simplifies Windows-to-Mac porting.

As for the Visual C++ 2.0 core package, it's an outstanding product. For new development work on Win32, Visual C++ and MFC create an unmatched combination. The only significant drawback is the chasm between 16-bit and 32-bit applications development imposed by the splitting of the two compilers and the abandonment of VBXes on 32-bit platforms. The environment is good enough that we wish we could use it to target 16-bit applications. ∎

Steve Apiki is a BYTE contributing editor and senior developer at Appropriate Solutions, Inc., a Peterborough, New Hampshire-based consulting firm specializing in multiplatform development. You can reach him on the Internet at apiki@apsol.com.

Desktop-To-Unix Network Integration?

Unix

Unix

Desktop PC

Desktop PC

Ours

Theirs

Take The Easy Way Out

If you're frustrated by the confusing maze of NFS-based desktop-to-Unix networking, we can show you an easy way out! PC-Interface 5.0 from Locus Computing is a simple, easy-to-use solution that adds PC server capabilities to your Unix System. Install PC-Interface 5.0 in minutes and suddenly you can easily share Windows, DOS, Macintosh and Unix files, applications and printers. No more complicated NFS mount tables! PC-Interface 5.0 screens look just like the Windows, DOS or Mac screens you're already comfortable with—*no need to learn Unix!*

Not only is PC-Interface 5.0 incredibly easy to use, it contains the most powerful, enterprise-wide feature set on the market. You get "gateway" access to any NFS, AFS or DFS system. You can integrate Novell NetWare, Unix and Windows for Workgroups networks. TCP/IP transport and WINSOCK for client/server applications is included. If you're a Unix user, you'll appreciate the consistency of standard Unix commands, tools and file access security. For even more power, Locus' PC-Interface Plus 2.0 also features enterprise-wide and global E-mail, advanced terminal emulation and FTP applications.

PC-Interface 5.0

Step Up To World-Class Network Integration

LOCUS
Computing Corporation

Take the easy way out TODAY! In the U.S.

Call 1-800-95LOCUS

Contact Locus on the Internet:
World Wide Web: http://www.locus.com
Internet e-mail: ussales@locus.com
                              eurosales@uk.locus.com (Locus UK)

Or phone or FAX the following Locus sales offices and distributors:

North America: 1-800-95LOCUS (US only), (310) 670-6500, (310) 670-2980 FAX
Europe, Middle East, Africa, Asia: LOCUS UK, (+44) 1442 236111, (+44) 1442 236453 FAX
Latin America, Asia/Pacific: (310) 337-5017, (310) 670-2980 FAX
Mexico: MPS Mayorista, (525) 325-0993, (525) 687-8163
Mexico: Premium International, (525) 525-1479, (525) 687-0425
Australia: MUA PTY, Ltd., (612) 419-5799, (612) 419-8445 FAX
Australia: Unetix Systems, (619) 474-1184, (619) 474-1034
New Zealand: The Great Escape Company Ltd., (649) 443-2421, (649) 443-2453 FAX

Circle 301 on Inquiry Card (RESELLERS: 302).

Intrusion Protection for Networks

CheckPoint's FireWall-1 protects your gateways from the growing threat of Internet security breaches

J. BRUCE DAWSON

Like most useful tools, the Internet is a double-edged sword. It's an open, flexible resource that lets you log in to systems around the world, transfer files, and execute programs. But once you're connected, it's also a way for other people out there to get into your system.

An increasing number of businesses have begun to address these Internet security issues by constructing firewalls, or physical boundaries, between private and public networks. While it's possible to build router-based firewalls using in-house development, a potentially more affordable and robust solution is to run a commercial firewall product on a dedicated server.

CheckPoint Software Technologies' FireWall-1 is among the first stand-alone firewall products. It is noteworthy in part because Sun Microsystems' SunSoft division chose it to work with the company's Solaris-based Internet servers. FireWall-1 runs on SPARC machines running SunOS 4.1.3 or higher and Solaris 2.3 or higher, as well as Intel 386-or-higher computers with Solaris 2.4.

Building the Wall

FireWall-1 provides various levels of security from wide open to screwed down tight. It is also very flexible, permitting you to selectively control access at the service and host level, along with providing voluminous details on attempted and granted accesses. Additionally, it lets you run scripts to control entire classes of access, such as packets coming in through a serial-interface gateway or through a Telnet connection.

In the current version, FireWall-1 works by monitoring the lowest level of the IP network: the individual packets. Accordingly, it works only with Internet, or other TCP/IP-based, networks. It will work with other protocols, such as Novell NetWare, if they are being tunneled through IP. (CheckPoint says FireWall-1 is actually protocol-independent and may

be ported to other protocols, such as IPX, in the future.)

By convention, all network hosts using the IP suite are allocated at least one IP address from the networking authority (which is either the domain administrator or the national Network Information Center). The upshot: If your system has an IP address and is connected to other networks, it is open to attack. (For that matter, your computer is also vulnerable to attack from others on your LAN.) Although most small sites gain some level of protection from enforcement of computer use policies, larger sites are more vulnerable because policies are harder to enforce.

Enter the enforcer. FireWall-1 provides both a GUI and a command-line interface that allow a site to monitor all packets passing over the LAN. Additionally, if a given computer is acting as a router or gateway (i.e., if it connects two or more physically separate networks), FireWall-1 can perform the following operations on classes of packets that the system administrator can specify: It can drop the packet, thus preventing it from going through the gateway; it can issue an alert; and it can return an error message after a failed attempt to send a packet through, essentially rejecting the packet. And you've always got the option to perform a logging action

Access Rules

FireWall-1 employs a set of user-defined rules to judge the admissibility of network packets. Each rule specifies what is to be done with a packet based on its origin, destination, and packet type. The latter are most likely to be UDP (User Datagram Protocol) or TCP packets, depending on the particular site; they may also be entering via FTP (file transfer protocol) or Telnet. If a rule applies to a packet, FireWall-1 takes the indicated action (accept, drop, or reject) and performs any logging associated with the rule.

There are five monitoring and logging options: "ring" a bell and post a message on the firewall console, send an E-mail message, log a message to a file, run a program (such as a user-written program that, for example, examines packets to determine which directory tree FTP is accessing), or do nothing.

FireWall-1 has numerous features that network administrators, in particular, will love. First, you have full control over rule sets, including the ability to keep multiple sets and institute them for various times of the day by using FireWall-1's command-line interface and Unix's `cron` facility. And, besides the powerful command line, FireWall-1 comes with a library for programmers and a GUI based

FireWall-1's GUI is the main control center, where the network administrator identifies networks, gateways, and servers, assigning the desired level of security and transaction logging to each.

on OpenView, Hewlett-Packard's network management software.

At this level of control, FireWall-1 functions as a high-level monitor of network traffic, reducing its own negative impact on system resources but still assuring system security. It remains important, however, to use each network application's own security mechanisms; for instance, the file transfer program that invokes the FTP still depends on Unix's own file security tools (which include owner, group, and world read/write/execute permission).

That said, FireWall-1's real strength lies in its ability to apply reporting, monitoring, alerting, and logging to network logical links with minimal impact on the host system.

Going to the Wall

To get a feeling for FireWall-1's capabilities, we set up two Ethernet IP LANs and connected them at a computer running the software. The firewall machine was a Sun Microsystems SparcStation. At various times, we also used a Dell 32sli notebook running Windows and FTP's PC/TCP; a Fintronics 486/66 running Linux, a Unix variant; a DECStation 3100 running Ultrix, Digital's Unix variant; or a Visual X-Term terminal. Due to configuration problems, we were unable to test SLIP and PPP connections.

The first LAN (IP address 192.1.1.0) was an internal network with FireWall-1 as its only connection to the outside world. The other LAN, 199.125.76.0, had a local direct connection to the Internet. Additionally, we made sure that at least one DOS-based PC was running FTP's PC/TCP to test FireWall-1's ability to work with non-Unix IP systems. (FireWall-1 did so without a problem.)

After completing the network configuration on every test system, we created some rules that would permit "fingering" (generating a list of currently logged-on users) only from systems on the 192.1.1.0 network but would drop fingers that were transmitted into the 192.1.1.0 network. The system worked as intended.

Next, we got a little devious and tried to controvert the security rules by fingering from a computer on 199.125.76.0 to one on 192.1.1.0, but FireWall-1 would have none of that. The firewall computer appeared to hang, which is an awkward way to prevent unauthorized access. We then switched FireWall's action option from drop (which has a stop sign for an icon) to reject (which has a wrong-way sign) and tried again. This time, instead of hanging, the machine

---

FireWall-1's Scripting Language

FireWall-1 will load rule scripts written in its FW-1 language, which follows the basic syntax *scope action conditions*. You can also use FW-1 at the command line.

For example, the following code checks inbound packets at all interfaces on my-host and lets any host send mail to my-host:

```
#include "std.def"
inbound all@my-host accept
   dst in my-  host, smtp;
```

---

running FireWall-1 issued a more friendly "Connection Refused" message.

In a test using FTP, we told FireWall-1 to log all access attempts so we could later verify that setting the rules to accept, reject, and drop the connections had worked. They did; the log showed that each connection attempt was handled appropriately. We set up other rules to permit FTP and RWHO connections (the latter lists all users on the LAN) but to reject all others (including pings, signals a computer sends out to see if it can reach the host). We verified that the log contained the desired information.

Some Shortcomings

Other tests included loading new rule sets with connections in progress—a process that revealed some minor shortcomings of the product. CheckPoint says loading rule sets is an atomic operation (i.e., one that must be performed entirely or not at all), so it is not supposed to be possible for packets to slip through during rule loading. We discovered while switching rule sets that supposedly protected networks became exposed if the firewall system was kept on while FireWall-1 wasn't running. There appeared to be no way to keep the gateway buttoned up in this situation. CheckPoint says the network administrator can configure the boot process to determine which modules and applications load first, thus preventing any packets from slipping in.

We also wanted to see what would happen if we changed rule sets with an FTP link in progress. We found that a ruleset change does not affect an FTP file transfer that is in progress, but it affects subsequent FTP connections. CheckPoint explains that this is so because with the

---

default rule set in effect, FireWall-1 looks only at the initial logical connection of the TCP port. To avoid this scenario, the network administrator can override the default by changing the "Enable Established TCP Connections" option in the GUI's "Control Properties" window, according to CheckPoint. Then, while an FTP session is in progress, if the rule set is changed to prohibit the FTP service from passing any packets at all, the user's FTP session will be interrupted. Exactly how that session is interrupted (either by hanging the particular connection or by sending an error message, as in the earlier example) depends on the new action specified by the administrator (e.g., a reject instruction might produce an error message, and a drop instruction would cause the connection to hang, and FTP would call time out).

The Final Checkpoint

Running FireWall-1 can be a good way to enforce a company's Internet access policies because it offers extensive monitoring and reporting facilities with controls for preventing anticipated intrusions. Most important, it permits flexibility in Internet access policies, allowing exclusion of anything from an individual packet to entire classes of services. Additionally, if an unanticipated security breach occurs, the logs generated by FireWall-1 might hold evidence that will help track down and prosecute the culprit.

In short, FireWall-1 is a powerful and flexible tool for combating the growing threat of network security breaches originating in the often anarchic world of the Internet. ∎

J. Bruce Dawson is a consultant with Virgin Software Ltd. in Manchester, New Hampshire, and has been a developer of low-level Unix, VMS, and DOS applications for 10 years. He can be reached on the Internet as jbd@virgin.mv.com or on BIX c/o "editors."

Business Objects Done Right

Next shows it means business with Enterprise Objects Framework, a data management infrastructure that moves a generation beyond existing tools

KEVIN SVEN BERG

he current trend in development tools is to tout visual construction environments, object frameworks, and increased data access capabilities (witness Powersoft's PowerBuilder, Borland International's Delphi, and Microsoft's Visual Basic). However, some visual environments share few of the benefits of true object orientation, and some object frameworks deliver little more than an underlying procedural orientation. Likewise, many database tools merely present the guise of visual manipulation. It is easy to construct whiz-bang demonstrations and prototypes with tools such as Next Computer's DBKit, but what does it take to construct real-world, corporatewide solutions in the business environment?

The short answer is reliable, industrial-strength database engines that are scalable across multiple networked platforms and easily modified. Next has created such a development tool in Enterprise Objects Framework, a coherent infrastructure of robust object systems for manipulating database information. EOF excels in its flexibility, letting developers factor, network, and distribute customized components. In addition, EOF is designed to minimize the impact of changing databases or data formats. Its simplified memory management and error handling make for more reliable programs. And ample visual tools make EOF easy to learn and maintain.

Working first with a beta version of EOF, I helped develop an EOF-compliant API for my company's compound document products. Although they are EOF-compliant, our current products do not directly depend on EOF.

The Framework

The software's framework design clearly distinguishes it from competing tools. While the packaged behavior of APIs limits a database library, a framework uses inheritance, protocols, and delegation to change a program's function. Object frameworks have become the de facto method of creating new applications, and specialized frameworks, such as Borland's OLE class library, are emerging as well.

Enterprise Objects Framework leverages the visual construction environment of the NextStep operating system's Interface Builder, integrating easily with other NextStep programs. Above, an EOController object (an instance of EOController, the heart of the interface layer) is wired to a column so that the DepartmentName attribute will be displayed. A second EOController (called Employees) has been arranged to create a master-detail relationship.

Taligent—the object-oriented operating system that IBM, Hewlett-Packard, and Apple are developing jointly—will compete closely with EOF as a more generic, framework-based development tool.

From a deployment standpoint, object framework design and implementation will determine how effectively you can leverage development efforts over the life of a product. You must look beyond the visual aspects of the current development tools and evaluate the underlying architecture. EOF's design borrows from the proven Smalltalk MVC (Model-View-Controller) paradigm, which cleanly separates data content from presentation with an intervening control layer. Adjoining components know just enough to communicate, so you can blindly exchange them like stereo components. The same paradigm has been applied at different levels of abstraction throughout the entire framework. Component factoring even extends into the data access layers, allowing more control over transaction management, interface and data buffering, and error delegation. All of this means that EOF is likely to meet real-world needs and adapt easily to new requirements.

The gestalt view of the architecture is intriguing. In essence, Next has created an intelligent data-flow conduit that can be interconnected, chained, and adapted to fit the terrain of an application, even a networked one. The framework becomes the infrastructure for the flow of data containers, which are called *enterprise objects*, derived from data models created using the included EOModeler application.

These objects contain key-value information that can map onto stored data at one end and into interface elements on the other. Access methods hide whether a value was extracted from a database record or derived from another source, such as a database join. Enterprise objects are the primary reason data format changes can be isolated from other parts of an application, reducing maintenance costs.

Building Blocks

The entire EOF is constructed from the new OpenStep Foundation Kit classes, fundamental building blocks that bring strong synergy to EOF components. The OpenStep specification is available via the WWW (World Wide Web) on the Internet at http://www.next.com, http://digifix.digifix.com, or ftp://ftp.next.com/pub/OpenStepSpec in either compressed PostScript or RTF format.　　　　*continued*

We were struck by the elegance and power of the Foundation classes compared to the tools that come with Next's NextStep operating system. Instead of writing code to parse a string into components or to read input from a file or to scan an input buffer for processing, you simply use a single-method call to the Foundation array, string, or scanner objects.

The most radical Foundation change involves memory allocation. The new auto-release memory policy eliminates the need to manage memory ownership between object relationships. While this simplifies normal coding, it also makes memory management of distributed objects completely transparent. This is where enterprise objects take a quantum leap by becoming autonomous data agents able to flow anywhere in a heterogeneous, networked enterprise framework, much like blood cells in an organism.

Using Enterprise Objects

Just how effective is the EOF architecture? We found it flexible enough to accommodate such unnatural experiments as replacing an SQL data source with a word processing engine. However, we found the EOModeler application less than hospitable to custom adapters and models, so we encoded them manually. Specifically, it was hard to debug custom code inside EOModeler. In theory, EOF layers can be made independent to the extent that vendor-specific information (e.g., SQL command variations) does not creep into portions of code or the model definition. In practice, it is easy to leverage SQL directly or to take advantage of specialized access at the adapter level, breaking modularity (Next even does this with a few examples). While this flexibility is a boon, EOF can sometimes give you just enough rope to hang yourself.

Furthermore, the auto-release memory policy, though it makes memory management nearly invisible, still presents challenges. It can be downright tedious to find mismatched retain/release pairs; imagine

debugging unpaired lock/unlock focus calls across numerous views in different programs. The suggested debugging methods provide too little control or information.

Also, with auto-release, memory performance tuning becomes a subtle art of balancing automatic versus intentional memory-pool reclamation. Fortunately, EOF already handles many memory issues efficiently and supports a broad degree of further tuning through fetch-on-demand and object uniquing strategies.

EOF relies on the first available components for OpenStep that descend from a new, incompatible root class. If you are integrating EOF into legacy code, you will face mixing objects from both new and old class hierarchies. Next has greatly smoothed the transition by using categories to implement shared protocols. However, you will still want to read the on-line manual, *Working with Interface Builder*, for advice on surviving in a "mixed" world. There, for instance, you'll learn about the need to use special archiving functions when storing objects from one hierarchy in another hierarchy.

Unfortunately, several useful EOF coding techniques are not mentioned in the *Developers Guide*. For instance, a top-level exception handler that retrieves error descriptions helps pinpoint problems much faster than error codes. Next should install this exception handler in version 1.1 as a debug feature for standard project templates. Another enhancement would be for Next to generate a description method for the enterprise object templates EOModeler creates.

Consider several important issues before adopting EOF. First, EOF is a NextStep-only solution. However, the number of supported platforms continues to grow (the list currently includes Motorola 680x0, Intel,

HP-PA, and SPARC), and EOF will support coming OpenStep environments, such as the Solaris implementation (see the textbox "What's Next for EOF"). If you're an independent software vendor, you may wish to consider whether your customer base has adopted EOF, because the required run-time libraries do not ship with NextStep. This makes adopting EOF an all-or-nothing commitment, which may discourage some customers. Finally, if you have a legacy application that relies on DBKit, be aware that Next regards EOF as the environment for future dvelopment, but the company will continue NextStep support for legacy applications.

If you do purchase EOF for personal use or custom development, note that EOF doesn't come with a database. This is not a huge drawback, as the product includes example code that uses a flat-file data source and inexpensive relational databases are available. We would like, however, to see B-tree support included in the base product as another efficient storage alternative.

The Object Isn't Objects

Next faces a challenge in marketing EOF, the product being neither fish nor fowl, but an infrastructure. It will be tempting to mimic vendors of other object-oriented tools and rely on whiz-bang demonstrations to market EOF. However, the real value of EOF lies in less tangible aspects, such as scalability, modularity, extensibility, and reliability—assets that are crucial to any corporate database solution.

EOF is a significant technology that warrants attention. With it, Next has inverted how we think about building database applications in exactly the same way object frameworks have changed the way we build other applications: outside-in versus inside-out (i.e., adding or modifying behavior versus calling libraries). The challenge for anyone evaluating EOF is to see beyond its interface wizardry into the essence of what it takes to run a business in the information age. ∎

Kevin Sven Berg is a Senior Software Engineer at Pages Software, Inc., a San Diego developer of NextStep-based compound document technology. He can be reached on the Internet at kevin@pages.com or on BIX c/o "editors."

One Slow Cow vs. One Fast Cat.

While they're grazing, ZEOS is blazing... and setting new records. In a December 6, 1994 comparison report, *PC Magazine* benchmarked 96 Pentium processor-based systems. The ZEOS Pantera Pentium-90 earned the highest Winstone 95 score of *all* machines tested, and the Pantera Pentium-66 easily passed both of Gateway's Pentium-90 PCs! If you're after performance (and who isn't?), then get a Pentium processor-based system. But, as the benchmarks prove, not all Pentium PCs deliver. The ZEOS Pantera will give you unbeatable performance and the highest quality features—all at an affordable price. And it's ready to ship *today!* For the fastest, hottest Pentium-based PC, call a ZEOS Systems Consultant now at 800-554-5226.

Winstone 95

Score	0	20	40	60	80	100	120	140	160
ZEOS Pantera Pentium-90									150
ZEOS Pantera Pentium-66									132
Gateway P5-90									113
Gateway P5-90 Family PC									109

WORST — BEST

Reprinted from PC Magazine, December 6, 1994

Fax Orders: 800-362-1205 or 612-362-1205. Phone Orders: Outside U.S. and Canada: 612-362-1212. Government: 800-245-2449. ZEOS Information Systems, Inc. GSA # GS00K94AGS5176. Purchase Orders, MasterCard, VISA, Am Ex, Discover, Z-Card, COD and affordable leasing programs.

Pantera As Tested:

Pentium-66	**$2745**
Pentium-90	**$3095**

PC Best — February 1995

- Genuine Intel® Pentium™ 66 and 90 processors
- 16MB RAM, 256K SRAM cache
- 1GB local bus IDE hard drive, 256K cache
- 4X CD-ROM drive and 3.5" 1.44MB floppy drive
- Diamond Stealth 64-bit PCI local bus SVGA color graphics card with 2MB VRAM upgrade
- 15" 1024 x 768 NI SVGA color monitor, .28mm dot pitch
- Two high-speed serial ports, one enhanced parallel port
- On-board Fast SCSI-2 option (Pentium-66) On-board PCI local bus Fast SCSI-2 and Ethernet options (Pentium-90)
- Six-bay desktop case (Pentium-66) Ten-bay vertical case (Pentium-90)
- 101-key space-saving keyboard
- Microsoft Mouse
- Microsoft DOS 6.2 and Windows for Workgroups 3.11
- Lotus SmartSuite

800-554-5226
24 Hours a Day
365 Days a Year

ZEOS INTERNATIONAL, LTD.

A playground for experts.
Training wheels for their friends.

THE NETWORKED PLANET

Traveling the Information Highway

A new exhibit at
The Computer Museum
computer_info@tcm.org / 617. 423. 6758 / www.net.org
300 Congress Street, Boston, Massachusetts 02210

Principal Sponsor: Sprint
Major Sponsors: Apple Computer, Hewlett Packard, Novell,
NYNEX, Stratus, S.W.I.F.T.

Radical Xbase Objects

CA-Visual Objects is a classy, object-oriented middleware package that opens a migration route for taking Clipper to Windows

RICK GREHAN

omputer Associates International's CA-Visual Objects is hard to classify. It is an application builder, an integrated development environment, and a programming language with a native-code compiler. It's also a migration path for bringing Clipper programs from DOS to Windows.

Perhaps CA-Visual Objects is best described as visual, object-oriented middleware, with *middleware* defined as database application generators that take much of the typing out of development. This significantly reduces development time, making CA-Visual Objects also a player in the growing RAD (rapid application development) market.

Nested Views

Working in CA-Visual Objects takes you into a three-tiered hierarchy of abstraction: application, module, and entity. These do not imply inheritance; rather, they serve as a mechanism for gathering and organizing the components of an application.

To get a quick idea of how the three elements fit together, imagine you have constructed a typical database application, say an inventory program. Fire up CA-Visual Objects, and the opening window shows you the "application view." Applications are represented by large buttons, each bearing a name and an icon that indicates the nature of the application: executable, library, or DLL (see the screen).

Double-click on your inventory application icon and another window opens. This window shows the modules in your inventory application. Its structure is much like the application view: named buttons decorated with icons. Modules are the gathering mechanism for the third and lowest level: the entity level.

Within modules reside entities, the "leaf" nodes of the imaginary tree. Entities are not grouping mechanisms as are applications and modules. Entities do the actual work: They are the menu defini-

The left subwindow is CA-Visual Objects' application view; the right subwindow is the module view (modules contain entities, the third and lowest level of the CA-Visual Objects hierarchy). Colored indicators on each button denote whether an entity is up to date (green), or modified (yellow) or contains an error from the last compilation (red).

tions, window definitions, classes, methods, field specifications, and so on.

Double-click on a module, and you open the entity browser window. This window looks markedly different from the application and module browsers. Each entity is represented by a horizontal button labeled with the entity's type and name. You open an entity by double-clicking on its button; it is this act that reveals CA-Visual Objects' real strength. For when you open an entity, you launch one of seven editors, chosen on the basis of the type of entity you've opened.

Click on a class or method entity, and you open the source code editor, which looks and operates like any typical text editor. Click on a window entity, and you launch the window editor. This is a visual window-builder, complete with floating toolbar that lets you drag and drop radio buttons, text boxes, and so on.

A central repository keeps track of dependencies between applications, modules, and entities. This architecture also makes recompilation faster (because changes can be localized at the entity level) and allows on-the-fly testing of code changes (because the repository keeps track of each application's components).

Noteworthy Entities

As mentioned above, all action takes place in entities, so some specific examples are worth exploring. A window entity is a window, but there are several kinds of windows in CA-Visual Objects. There are data windows, dialog windows, and shell windows. Of these, the data window is the most interesting because it connects to a data server and thereby becomes *data aware*. This means a data window "knows" how to display and to manipulate data from a database. You can display data in "form" mode—where each field is a separate, labeled edit box—or in "browse" mode, which presents the familiar grid of records arranged in rows, columns, and fields.

A server entity acts as an interface into a database—a view, if you will. In the server editor, you specify such things as what database file the server accesses, whether it is linked to another table in master-detail fashion, and what index files to use. The server editor also provides a list of the fields in the table.

Inside most CA-Visual Objects editors, you will find a floating, always-on-top window called the properties window. It lists the properties of whatever element currently has the focus in the active window. Open the server window, for example, and the properties window includes a list of such items as whether the file is

shared, whether it's read-only, or what the database driver is.

Clipper Port

If you're a Clipper programmer, you are eyeing CA-Visual Objects as a way to get your applications into Windows. The CA-Visual Objects language (called VO by Computer Associates engineers) is touted as a superset of CA Clipper. There are some differences; one of the manuals explains porting from DOS Clipper to CA-Visual Objects and includes an appendix on known incompatibilities.

Still, if you want to turn your whole under-65-KB Clipper program into a single, monolithic method—frightful as that sounds—you can do it. You'll use the DOS terminal emulation window class, and your application will end up looking just like a DOS program running under Windows. More likely, you'll end up migrating pieces of your application into the CA-Visual Objects paradigm. Luckily, your core code can remain largely unchanged.

The Object of Desire

We have seen similar products before: KnowledgeWare's ObjectView, Gupta's SQLWindows, and Powersoft's Power-Builder, to name three prominent ones. Few, with the possible exception of Power-Builder, are as well integrated as CA-Visual Objects. In some of the other packages, when you open, say, the equivalent of a window editor, you get the feeling you've been handed off to another program. Not so in CA-Visual Objects; you know you're just down in another one of the editors.

There's a lot more to CA-Visual Objects. When you build a new application, the system automatically gathers enough entities to make a working framework, eliminating piles of drudgework coding. Application windows come not only with prebuilt (but modifiable) menus but also floating toolbars. CA-Visual Objects supports Xbase, FoxPro, ODBC (Open Database Connectivity), and a number of SQL database drivers. And the interface grants full access to the Windows API and lets you call and create DLLs.

The list goes on. It all adds up to quite a package. CA-Visual Objects is an efficient, fully integrated database development environment. And if you're a Clipper developer, it's the best game in town. ∎

Rick Grehan is technical director of the BYTE Lab. You can reach him on the Internet or BIX at rick_g@bix.com.

Circle 94 on Inquiry Card.

Server with a Slot

With two PC Card slots, Microplex's multiprotocol print server can simultaneously support Ethernet and Token Ring LANs

BEN SMITH

With two PCMCIA slots for network cards (rear of unit, inset), Microplex's M204 print server can connect to two different networks at the same time. These can be a combination of thin Ethernet, UTP (unshielded twisted pair) Ethernet, or Token Ring. PCMCIA network cards are optional. All printer ports (two parallel and two serial) are bidirectional. Not shown is the unit's power transformer.

In the features-for-the-price war among stand-alone print servers, Microplex's M204 is winning, and it has little to do with the unit's PCMCIA network adapter design. While the two PC Card slots were what first caught our attention, the M204's other capabilities impressed us more (see "Features: More than PC Card Slots"). The $595 M204 is a small (8.5- by 6.2- by 1.3-inch) box with four printer ports (two parallel, two serial) and two PC Card slots for Ethernet and Token Ring network adapters.

As a multiprotocol print server, the M204 can serve your whole network, supporting PCs, Macs, and most Unix systems. As a stand-alone print server, the M204 gives you flexibility: You can connect nonnetwork printers directly to your network without an intermediate host computer, and in the location most convenient for sharing. Compared to most printers with a built-in network interface or network card option, a stand-alone print server like the M204 delivers more features and delays obsolescence. When upgrading printers, you don't have to buy a new network interface. And with its PC Card slots, the M204 can also upgrade to future network interfaces.

Strong on Features

The M204 has other slick features. Take, for example, the M204's ability to balance the network printing load by redirecting print jobs. You can assign any M204 printer port, serial or parallel, to one of eight "destinations" (logical queues). You can also assign any logical queue to an alternate queue. If the primary queue is busy, the Microplex server can automatically redirect print jobs to the alternate queue. Compared to directing jobs through a sophisticated print-spooling program (often single protocol) running on something like a Unix server, it's an easy and inexpensive way to double the capability of a print station.

The M204 can also handle bidirectional communications with attached printers through any one of several job status reporting methods. You can send a print job to a PostScript printer and have return messages sent to a console or even to an E-mail address. While some PC-based print drivers know what to do with messages returned from a bidirectional printer port, the M204 gives you the same level of management with network printers.

Despite the preponderance of parallel-port printers, the M204's two serial ports are quite useful. Hook up a terminal or a laptop with a communications program, and you can log in to the M204 to read or set any of its many configuration options. The serial ports can also serve as write-back paths for printers that support it. Most surprising, you can use an M204 serial port as a network terminal server to telnet into any address on your net.

Also unusual, the M204 lets you establish an encrypted data link with the host so that snoopy workstations and network sniffers can't see the contents of files as they move across the network. Not as secure as an RSA link, Microplex's scrambling technique does offer some level of protection and certainly privacy.

The M204's most distinguishing feature is its PCMCIA interface. The two PC Card slots aren't for flash memory (though the M204 has internal flash memory for upgradability); they're for PC Card network adapters, a first for stand-alone print servers. Not only does the M204 automatically detect network type and connect to two networks, the firmware supports simultaneous printing from Ethernet and Token Ring LANs. With two networks connected, the M204 can also act as a low-cost router. Future plans include using the PC Card bus for other devices as well, such as SLIP connections in or fax/modem connections out.

While the PC Card interface is flexible and convenient, it adds to the total cost of the package because the network interface cards are options. Microplex sells two $295 Ethernet cards (10Base-2 or 10Base-T) and a $495 Token Ring card. You can use cards only from Microplex, because each card requires specific M204 firmware support. It's not a notebook that can load Card and Socket Services from some PCMCIA card vendor. In this respect, the PC Card slots are little better than the proprietary interface upgrades provided in a print server like Milan Technology's FastPort 3200X.

Shell Collection

A product as complex as a multiprotocol print server requires good management software and documentation. Microplex provides both. The Macintosh and Novell operating systems already have all the management tools you need to get started. The Mac automatically sees the M204 printer ports on the network. Simply select the printer port from the Chooser menu.

Similarly, the NetWare supervisor uses the same Novell tools as would be used for any other network printer. Microplex provides complete installation instructions for NetWare, with the M204 appearing both as a NetWare print server and a network printer. Microplex includes an MS-DOS program

NPSH for monitoring the M204 and, with password protection, setting the configuration directly from a NetWare workstation.

For connecting the M204 through TCP/IP to Unix workstations or servers, your options are legion. The installation would be daunting were it not for the utilities and awe-inspiring management scripts that Microplex provides. You get compiled executable utilities for SCO (The Santa Cruz Operation) Unix 3.x, Silicon Graphics IRIX 4.0.5, and Sun Microsystems OS 4.1.1, as well as the source code to these utilities.

The Unix utilities include a method of sending data directly to any of the M204 ports, a newline to carriage-return-plus-newline filter, a scrambling filter for en-crypting the data going over the network, a banner text generator, a utility for wrapping ASCII text into Post-Script, and a daemon for handling return communications from the print server.

The most impressive piece of work is a 27,000-line shell script, `npconfig`, and a 7000-line subset of that script, `ezsetup`. Any Unix system can run the configuration shell scripts without special binaries, a compiler, or anything other than the utilities found on every Unix system. The disadvantage: Scripts this large are difficult to write and debug. Unlike the C programming language, or even Perl, shell scripts lack development environments and debugging tools.

Worse, shell scripts are slow. Microplex's shell scripts not only help configure the M204 but also help set up the print-spooler configuration on the Unix host—a nasty job on some Unix systems. It is often simpler and faster to configure the M204 through a telnet session.

The 200-page manual, well organized and well written, supplies a good index and a glossary specific to Microplex terminology, which are helpful because some concepts are a little difficult to grasp at first, and all the information that you need is there, with a few exceptions. Lacking, for example, are guidelines for troubleshooting NetWare installations.

The Shakedown

The M204 is impressive, but Microplex has a few design details to work out. Besides bugs in shell scripts (which are easy for any Unix administrator to fix), we found that setting up the alternate device paths is less than intuitive, some Macintosh error messages were inaccurate, and elements of the Unix installation were not as flexible as they should have been.

If Microplex put as much effort into developing a good socket-connected management client as it did into writing the gargantuan shell scripts, the TCP/IP configuration and control would not only be faster but far more flexible. Nonetheless, the M204 is a great performer and has tremendous flexibility for its price. There's also the promise that Microtest will upgrade the M204 firmware to support future PC Card network interfaces. ∎

Ben Smith is a Unix and IP networking consultant and a former BYTE testing editor. You can reach him on the Internet at ben@ronin.com or on BIX as "bensmith."

Project Management for Windows

Four midrange power tools that can help you marshall complicated projects to completion

SCOTT HIGGS

A s organizations merge, split, and downsize, managers face increasing pressure to improve efficiency and to document actions taken toward that goal. And they are turning to project management software for help. Project management programs provide a single interface to track and manage all the people, resources, and decisions employed to reach a goal. While using time-honored models such as Gantt and PERT charts, the current generation of project managers also uses newer E-mail and workgroup concepts to involve groups of people in project tracking.

NSTL tested the top four "midrange" project managers: Computer Associates' CA-SuperProject 3.0 for Windows, Microsoft Project for Windows 4.0, Scitor's Project Scheduler 6 for Windows 1.5, and Symantec's Time Line 6.1 for Windows. All four programs carry price tags under $700, target nonspecialists, and support workgroups and enterprise connectivity. NSTL limited its review to products that can model complex task relationships, provide full reporting and charting capabilities, and manage resources with automatic leveling.

Keeping on Time

All the tested programs include fairly complete tools for scheduling tasks. You can specify durations down to the minute, place constraints on start and finish dates, and calculate schedules by applying either "as soon as possible" or "as late as possible" models. Tasks and projects can extend over decades, if desired, and the programs readily handle all standard task-precedence relationships, calculating float (the length of time an activity can be delayed without disrupting the project) and applying resource calendars as they do so.

A handful of special capabilities distinguish the programs. Project Scheduler 6 lets you assign unique calendars by activity, as well as the standard resource calendars. Time Line for Windows offers more flexibility in scheduling unusual work shifts than do the

The horizontal bars of Gantt charts (right of screen) are standard fare in project management software. However, Microsoft Project for Windows (shown here), was singled out among the four tested programs for ease of learning and usability.

other products. CA-SuperProject for Windows and Project Scheduler 6 let you enter best, worst, and most likely estimates for task durations and then use these estimates to calculate probabilities for project dates.

Tracking and Reporting

Some managers use project management software simply to draft a schedule and post a chart on the wall. But increasingly, managers find that tracking and entering progress data can be worth the time. The extra effort allows printing up-to-date status reports and graphs, identifies problems before they snowball, and creates an archive of data that can help improve the accuracy of future schedules.

All four programs tested encourage you to continue using the software throughout the project. They will calculate remaining durations and costs, compare budgeted and actual costs, and generate profiles of resource usage. With the exception of CA-SuperProject for Windows (which offers only a global update), all make it relatively painless to update progress on a number of selected tasks all at once.

CA-SuperProject for Windows and Microsoft Project for Windows offer unusual tracking flexibility by letting you

OVERVIEW

NSTL RATING		PERFORMANCE	QUALITY	VERSATILITY	EASE OF LEARNING	EASE OF USE	PRICE
★★★★	Microsoft Project 4.0	▲	■	▲	▲	▲	$469*
★★★★	Project Scheduler 6	■	▲	▲	▲	▲	$695
★★★	CA-SuperProject 3.0	▲	▲	▲	■	■	$649
★★★	Time Line 6.1	▼	■	▲	■	▲	$500

* Estimated retail price

KEY

★★★★★	Outstanding
★★★★	Excellent
★★★	Average
★★	Below average
★	Poor
▲	Good
■	Fair
▼	Unacceptable

set up to six baseline schedules per project. This capability enables more efficient what-if analysis and provides alternative views of progress. The other programs maintain only one baseline per project.

All the products support extensive sorting and filtering options on reports. With the exception of CA-SuperProject for Windows, all also include a generous assortment of formatting options, including extensive font control, user-specified reduction and enlargement, and output forced to fit a page. Time Line for Windows and Microsoft Project for Windows provide handy calendar reports that display task assignments and progress in a familiar format.

Resource Management

CA-SuperProject for Windows delivers by far the most complete set of functions for modeling real-world resources. It allows you to define material resources, vary resource capacities and costs, factor overhead costs, and assign overtime based on task priority. Along with Project Scheduler 6, CA-SuperProject for Windows stands well ahead of the rest of the class with its ability to use resource leveling (automatic reallocation of resources to better balance the workload) to optimize a schedule by splitting tasks.

Of all the specialized functions mentioned above, Microsoft Project for Windows offers only overhead costing. For most other procedures, you must resort to a workaround, such as an external spreadsheet program for calculations. Project provides the tools for minimal resource management, but it has few refinements for handling complex resource situations.

The other two programs fall between these extremes. Time Line for Windows is a bit stronger in letting you set variable resource capacities, while Project Scheduler 6 permits nonuniform resource use and employs a more effective leveling algorithm.

Across the Enterprise

Vendors and managers agree that workgroup functions hold the key to the future: Almost

Highlights

	Strengths	Limitations
CA-SuperProject 3.0 for Windows	Fastest program overall	No ODBC, OLE support
	Summary PERT	Steepest learning curve
	Sophisticated resource leveling	Less flexible chart presentation
Microsoft Project for Windows 4.0	Customizable toolbar	No task-splitting option
	Flexible reporting, calendaring	Simplistic resource-leveling algorithm
	Easiest to learn and use	Weak tools for resource management
Project Scheduler 6 for Windows 1.5	Sophisticated resource leveling	No OLE support
		No macro debugging tools
	Unlimited undo	No font mixing in PERT charts
	ODBC support	
Time Line 6.1 for Windows	OLE and ODBC support Customizable toolbar	Sluggish performance
	Summary PERT	Simplistic resource-leveling algorithm
		No simultaneous task editing

Primavera Takes SureTrak in Two Directions

Primavera Systems, Inc. (Bala Cynwyd, PA; (800) 872-5457, (801) 973-1300), provides some of the industry's most popular software for professional project managers in manufacturing, construction, engineering, and a wide range of other fields. Their flagship product, P3 (Primavera Project Planner), comes in versions for DOS and Windows, both optimized for large-scale, complex projects. At $4000 per copy, P3 addresses the needs of specialists more than the scheduling goals of average managers. SureTrak Project Manager for Windows is Primavera's first entry into the midrange Windows market, priced to compete aggressively with the other products reviewed here.

SureTrak has capabilities similar to those of the four project managers we tested: fully customizable Gantt charts, split-screen views, automatic resource leveling with priorities, and cost and schedule variance reports. Primavera loaded many sophisticated resource-tracking and reporting functions into SureTrak to encourage managers to use the program throughout the entire life of a project. SureTrak also offers revenue fields as well as cost fields to accommodate resources whose costs can be recovered. Another unusual option is a function that automatically "spotlights" all tasks that should have been worked on in a specified time period. This lets managers focus quickly on key activities and assess progress.

For many users, SureTrak's major appeal is its tight integration with P3. Both programs can read and write data files in P3 format, minimizing the need for translation. Primavera encourages users to think of a "concentric" model for enterprise-wide project management, with all data in the same database and

different users employing different tools for access and editing. This model proposes that individual project managers work with SureTrak and that higher-level managers view and control multiple projects with P3. Managers with specialized needs can use add-in products (each costing from $2500 to $5000) with modules to support risk analysis, probability, and tight scheduling of intense, short-duration projects. For connectivity outside the Primavera family, SureTrak is also compliant with ODBC (Open Database Connectivity), allowing direct access to project data in other ODBC applications.

Based on a brief trial with the program, SureTrak was a strong product that should be competitive with the four tested programs. It has greater flexibility than most in resource management, and it offers a clear upgrade path for people who need sophisticated modeling.

GenIE
The power of intelligent Data Modeling

You've heard about CASE before, Computer Aided Software Engineering. It's supposed to help you **build better, more precise, more reliable, and ultimately more maintainable software,** but often it just doesn't seem to work out that way.

Perhaps you've tried CASE, but found it got in your way or made you less productive. Perhaps you're forced to use CASE for documentation purposes, but it doesn't integrate well into your development environment. Or perhaps you've personally mastered the benefits of CASE but want a less expensive way of having your entire team using CASE productively.

GenIE is designed for people who want more success, and more professional results from their existing development environments. It is for people who desire to become more effective analysts and want more innovative development teams. It is also intended for people who are struggling with providing **precise, understandable, professional system design plans, and finished system documentation.**

Entity Relations
Takes the pain out of database definition.

Function Hierarchy
Helps you cope with scope.

Data Flow
Show your system works.

Matrix Modeling
Maps your data usage.

Unique CUE-CARD Learning System

No expensive training required.
GenIE comes with our unique *Cue-Card* on-line tutor system. This is much more than help, because Cue-Cards not only tell you what you can do, but actually show you how to do it. With this system you can start using GenIE right out of the box, and learn as you go.

GenIE is Industrial Strength CASE.
Those familiar with CASE will appreciate the following specifications:
- Single repository for all modeling objects and diagrams.
- Access repository via SQL/DDE link from any Windows application.
- User configurable reporting.
- Print frames support WYSIWYG and multi-page formats.
- Configurable diagraming notations gives consistency with major methodologies.
  - Selectable validations ensure database correctness.
  - Multi-user client-server technology with roll back recovery.
  - Automatic SQL, DDL and xBase data, index, and memo generation includes ANSI, Gupta, Informix, Clipper, FoxPro, DB2, Access, Oracle, and SYBASE.
  - Portable relational repository: UNIX, Novell NLM, OS/2, NT.
  - Non-prescriptive approach does not force any methodology.

GenIE has simplified CASE.
This product simplifies the language of CASE—making it more accessible to everyone. Now you don't have to be trained in CASE methodologies to take advantage of these valuable tools and techniques.

"Previous tools of this sort have been an overhead to our DP development. This is much easier to use and offers real improvements in productivity."
-John Little, DP Development Director, Cambridgeshire County Council

For a complete business process reengineering tool, try our **Modeler v4.0** with the revolutionary *Process Simulator.*

Call Now
800-893-8398
FAX: 617-894-0418
Call or FAX for more information.
30 day test drive version available.

Outside the U.S. contact:
CASEwise, Inc.
Swiss Cottage House, 3rd Floor
9-13 Swiss Terrace, London
England, NW6 4RR
TEL: +44 (0) 171 722 4000
FAX: +44 (0) 171 722 4004

GenIE integrates into your development environment under MS Windows™ as an intuitive design & analysis component.

Dealer inquiries welcome.

Circle 281 on Inquiry Card.

nobody works on projects in isolation. Each vendor takes a unique approach to supporting the enhanced connectivity required for project management.

Time Line for Windows emphasizes connectivity to other corporate data in a client/server model. It features a built-in SQL database engine, permitting you to integrate Time Line data with other core business applications for reporting, updating, and analyzing information. This model assumes that a company has a central data repository and that Time Line for Windows will be just one of many specialized tools for working with subsets of that data.

Project Scheduler 6 supports many of the same benefits as Time Line by virtue of its strong ODBC (Open Database Connectivity) support. Microsoft Project for Windows and CA-SuperProject for Windows offer good import and export capabilities but are less suited to sharing enterprise data. Time Line for Windows and Microsoft Project for Windows provide the most complete OLE support, and all products except Project Scheduler 6 include DDE for linking with other Windows applications.

All these programs can establish central resource pools for multiple projects and generate reports and charts that represent an overview of several projects. CA-SuperProject for Windows makes some of these tasks relatively awkward, requiring you to "combine" projects to generate overview charts. Microsoft Project for Windows can establish task dependencies between projects, but you must set up DDE links rather than use an internal program function (as offered by other vendors) to handle this procedure. Time Line for Windows and Project Scheduler 6 manage multiple project models with fewer workarounds.

Communication problems among team members can often thwart a project's success. Therefore, Microsoft Project for Windows includes unique tools to expedite communication and to minimize repetitive data entry. Thanks to tight integration with Microsoft Mail (and other MAPI-compliant mail systems), managers can plan a schedule, click a button, and have automatic mail messages generated to all the people who are assigned to tasks. Built-in messages (including check boxes and a notes field for recipients) offer initial confirmation of assignment, updates (when schedules change), and progress tracking. When reviewing replies, the project manager can incorporate the information received at the click of a button, automatically updating the master project file to reflect the status reports received. This time-saving feature dramatically improves the usability of project management software beyond that of the initial planning stages.

Mid-level project management software still lacks an effective architecture for managing simultaneous access by multiple users of corporate data. Currently, none of these

PROJECT MANAGEMENT FEATURES

	CA-SuperP	MS Project	P Scheduler 6	Time Line
PROJECT CALENDAR				
Maximum size (years)	179	65	100	120
Latest calendar date	2129	2049	2079	2099
Breaks out calendar by activities	○	○	●	○
Breaks out calendar by resources	●	●	●	●
Multiple master calendars (for resource groups)	●	●	●	○
SCHEDULING CONSTRAINTS				
Start no earlier than ①	●	●	●	●
Start no later than ①	●	●	②	●
Finish no earlier than	●	●	②	○
Finish no later than	●	●	●	○
Fixed (mandatory) date activities	●	●	③	●
Calculates ASAP/ALAP scheduling	●	●	●	●
TASK DURATIONS AND RELATIONSHIPS				
Durations in minutes	●	●	●	●
Durations in hours	●	●	●	●
Best/worst/most-likely estimates	●	○	●	○
Calculates total float	●	●	●	●
Specifies duration in elapsed days	●	●	○	●
RESOURCE INFORMATION				
Maximum resources per project	U	9999	10,000	U
Maximum resources per task	U	100	10,000	U
Sets resource capacities	●	●	●	●
Normal and maximum resource capacities	●	○	○	●
Variable resource capacities	●	○	○	●
Variable resource cost	●	○	●	●
Overtime resource costs	●	●	●	●
LEVELING				
Continuous resource leveling	●	●	○	●
User-specified resources to level	●	●	●	●
User-specified tasks/dates to level	●	●	●	○
Level material resources	●	○	○	○
Sets minimum resource allocations	○	○	○	○
TRACKING FEATURES ④				
Maximum baseline (target) schedules	6	6	1	1
Remaining durations for activities	●	●	●	●
Show percentage complete	●	●	●	●
Actual cost of activity	●	●	●	●
Summary update of project	●	●	●	●
Summary update of selected activities	○	●	⑤	●
Globally update percent complete	○	●	⑤	●
Budgeted resource amount	●	⑥	⑥	⑥
Budgeted resource cost	●	⑥	●	●

● = yes; ○ = no; U = unlimited
(Note: Features that are not integral to the main program are marked "○")
① Prevents schedule calculations from overriding critical dates
② Sets duration and finish/start no later/earlier than date
③ Simulated by setting start date constraints and entering a duration
④ In most cases, the user enters information and the program calculates results
⑤ Can update selected tasks to a user-specified date
⑥ User must enter figure

four project managers supports the record-level locking necessary to allow multiple users to edit different portions of a file at the same time.

Pretty Pictures

Gantt charts (the screen on page 185 shows an example from Microsoft Project for Windows) form the backbone of basic pro-ject management. Accordingly, most programs bring up a Gantt view by default and give users a fair amount of control over formatting the charts. The same basic information and look characterize all of these programs' Gantt charts, making them roughly comparable in terms of basic analysis. Variations in this category primarily reflect choices in time-scale layout, legend, and general formatting.

Project Scheduler 6 and Microsoft Project for Windows generate the most attractive Gantt charts, with clear time scales and a high-contrast layout. These two programs create the kinds of charts that will impress clients. CA-SuperProject for Windows and Time Line for Windows proffer respectable, readable charts that serve well for in-house purposes.

PERT charts exhibit much more variation, as each program takes a different approach to placing, spacing, and connecting the many nodes that represent project tasks. We found that none succeeds in the dual goals of clarity and attractive presentation. CA-SuperProject for Windows comes closest in the former category, with connecting lines that make task relationships relatively easy to follow in a layout only an engineer could love. Microsoft Project for Windows makes a much prettier picture but generates a chart full of links that are difficult to decipher. Time Line's charts are fairly clear but quite unattractive, and Project Scheduler 6 generates attractive nodes tangled in a cobweb of lines.

Roll Up

Managers who send proposals to clients or to top management must roll up selected project data into custom-formatted charts and reports. Just as advances in spreadsheet formatting and charting have raised standards for general data presentation, charting enhancements among project management programs have raised expectations for customized output.

Microsoft Project sets the pace in this area, with more variations offered (including free-floating text and graphics) than in any other program. Project Scheduler 6 provides strong graphical analysis, while Time Line offers a wide range of text-formatting options. CA-SuperProject for Windows brings up the rear, delivering the basics but imposing awkward limitations on page and text formatting.

Quality of schedule calculations seldom receives much attention in reviews. Differences become apparent only when applied to complex projects—exactly the circumstances faced in offices every day. NSTL created such a project, assigned resources,

PROJECT MANAGEMENT FEATURES

	CA-SuperP	MS Project	P Scheduler 6	Time Line
GANTT CHARTS				
Total float	●	●	●	●
Free float	●	●	●	●
Time scale in minutes	○	●	●	○
Time scale in hours	●	●	●	○
Scale by variable periods, not calendar dates	●	●	●	○
User-selectable on-screen chart size	●	●	●	●
Shows task-dependency lines on Gantt	●	●	●	●
Activities grouping in summary bars	●	●	●	●
Can select and mix fonts or styles	●	●	●	●
PERT CHARTS				
Zoom feature	○	●	●	●
Activity named directly on node	●	●	●	●
User-positioned tasks	●	●	●	●
User-positioned arrows	○	○	○	○
Automatic, time-scaled task placement	●	○	○	●
Displays lag times between activities	○	●	●	○
Precedence relationships	●	①	●	○
Summary PERT	●	○	○	●
RESOURCE REPORTING				
Resource report by activity	●	●	●	●
Budgeted vs. actual variance	●	●	●	●
Report in wall-calendar format	○	●	○	●
Resource profile as bar chart	●	●	●	●
Resource profile as stacked bar/multiline chart	●	●	●	○
Resource profile as line chart	●	●	●	●
Calendar view by resource	○	●	○	●
Time scale in hours	○	●	●	●
Cumulative resource profile	●	●	●	●
Resource capacity report	●	●	●	●
NETWORK CAPABILITIES				
Network (multiuser) version available	●	●	●	●
Locks by record (not by file)	○	○	○	○
Read-only file lock	●	●	●	●
Identifies lock holder	○	○	○	○
Maintains personal settings	●	●	●	●
Password protection	●	●	●	●
Passwords for different access levels	●	●	●	●
SUPPORT FOR MULTIPLE PROJECTS AND WORKGROUPS				
Combines multiple projects	●	●	●	●
Links multiple projects	●	●	●	●
Max. number of projects open simultaneously	②	80	②	U
Inter-project task dependencies	●	③	●	●
Maximum number of projects linked at once	④	②	②	U
Automatic update of master project	●	●	●	●
Central resource data for multiple projects	●	●	●	●

●=yes; ○=no; U = unlimited
(Note: Features that are not integral to the main program are marked "○")
① In special, less detailed PERT chart
② Limited by available RAM
③ Available via Windows DDE/OLE
④ Varies from 15 to 255, depending on DOS version

WHAT KIND OF EVIL-MINDED SLIMEBALL REPORTS THEIR SMALL BUSINESS COLLEAGUES FOR ILLEGALLY COPYING SOFTWARE?

When people call the SPA to turn in someone for illegally copying software, they do so because they know that when dishonest people steal software, the rest of us end up paying for it. Copying software is a crime. To clean up your software, or report someone who should, call 1-800-388-7478.

SOFTWARE PUBLISHERS ASSOCIATION

DON'T COPY THAT FLOPPY.™

CA-SuperProject 4.0: A Welcome Upgrade

C omputer Associates International was expected to announce a major upgrade to CA-SuperProject for Windows by early this year. Although this upgrade, version 4.0, was not available during the testing period, we took a close look at a prerelease version of the application and noted significant improvements in the program's usability and versatility.

The most visible improvements are in the user interface. Data entry forms now work in split views, and they can change their content to reflect user selections in other windows. This change allows more control over the work environment and greatly speeds editing. New tabbed dialog boxes make a wider range of options available as well as reduce the need to jump back and forth between menus.

Other improvements include spell checking and multilevel undo. Stronger mouse support allows more flexibility in graphical editing of charts and more control over task links during clicking and dragging. New views and improved charts make it easier to track projects. Gantt charts now have more flexible formatting, and histograms chart a wider range of fields. In addition, CA-SuperProject now tracks three baselines to facilitate comparisons and what-if analysis.

Managers who coordinate multiple proj-

ects will appreciate the new combined project view, which provides combined, editable views, all while maintaining each project as a separate entity. You can create task links between projects with the click of a mouse, and all linked projects recalculate simultaneously, as needed. The program now supports data exchange via ODBC (Open Database Connectivity) and as an OLE 2.0 client.

Overall, CA-SuperProject 4.0 for Windows addresses many of the usability concerns that frustrated NSTL's testers. New functions bolster this product's claim to offer the most-powerful features in this midrange market segment.

---

and determined a standard project end date within each program. Then, after changing one resource assignment, our technicians had each program recalculate and opti-

mize the schedule using resource leveling.

CA-SuperProject for Windows and Project Scheduler 6 came through this test with flying colors. These programs analyze an entire project, reshuffle task sequences, and propose new schedules that add only two days of delay to the project overall. Microsoft Project and Time Line fail to find any of the available shortcuts, propose lengthy periods of relative inactivity, and delay the whole project by 12 days.

No vendor algorithm can possibly return optimum results in all situations, but Project Scheduler 6 and CA-SuperProject for Windows present suggestions that require far less manual editing than those of Time Line for Windows and Microsoft Project for Windows.

Performance Matters
Project Scheduler 6 will irritate managers who are eager to see their printouts. The program takes over 3 minutes to finish printing a Gantt chart and more than twice as long as the next-slowest program to print a PERT chart (over 6 minutes total). All other program functions are fast, and even when printing, you'll regain control of your system in a few seconds.

Time Line for Windows records its best times in printing, and lags well behind the other programs in almost every other measure. The overhead of its SQL database makes file manipulations relatively slow. More important, routine editing changes that require schedule recalculation tie up the system over four times as long as the next-slowest program. With Time Line,

you'll wait 13 seconds for a recalculation that Microsoft Project and Project Scheduler perform in 1 second or less.

Seeing It Through
The programs reviewed here require much less specialized training than those of a previous generation. In the near future, planning schedules, assigning resources, and tracking progress with project management software may be as fundamental to management as using a spreadsheet has become for budgeting.

To choose the most appropriate software for your circumstances, consider the relative importance of usability, specialized project management functions, output quality, and workgroup capabilities. NSTL weighed these and many other variables in calculating our overall rating, based on our assessment of a typical user's needs. In the end, Microsoft Project for Windows and Project Scheduler 6 for Windows both earned NSTL's second-highest possible recommendation—four stars for excellent. ■

Not One More Damn Line of Code. Ever.

See How LAYOUT Lets You Build Real, Heavy-Duty Programs Without Writing a Single Line of Code. Free.

PC Week called **Layout** a "sure thing." We call it a revolution. With over 200,000 users, tons of add-ons, and widespread third-party support, Layout is the only tool that lets you build DOS or Windows programs by manipulating objects on screen — without writing code. Not just simple programs, but real, heavy-duty, mission-critical applications.

The True Power of Objects

Layout is truly object-oriented, both in the programs it creates, **and** in how you use it. You start out by arranging objects in a simple diagram, and then add more objects as the program grows, or create new objects by combining existing ones. You can even run your program as you're building it. Data-entry, database, and report formats are all visually designed on-screen.

What Layout Delivers

When you're done, Layout creates real .EXE files, or well-structured and efficient C/C++, Pascal, or BASIC programs. You can even create new objects right in Layout, or even re-use existing source code. Layout supports DOS and Windows, with NT and OS/2 coming soon, and applications written on any of these platforms are automatically portable to

the others — including Windows 95 (Chicago).

Visual Power, Incredible Performance

The programs Layout creates are completely graphical, even under DOS, and fully support OLE 2.0, DDE, 3D buttons, hypertext links, messaging, creating and using DLLs, and much more. Layout even supports pictures as a data-type! Layout creates very efficient programs — they're fast and compact. No 150K "Hello World" programs come out of Layout: it doesn't just spit out pre-canned code like other so-called high-level tools. And now,

Objects, Inc. is offering a free Windows95 Toolkit with Layout so you can start building tomorrow's programs today!

Layout: the Ultimate Tool

Layout delivers the future for $299.95, including everything you need to build programs; free, unlimited, technical support; and superb documentation. Call today, join the Layout revolution, and never see code again. Ever.

Free!
LAYOUT APPLICATIONS SAMPLER
☎ **800-424-6644**

Objects, Inc.

HANDS-ON TESTING

30 NO-COMPROMISE NOTEBOOKS

We test top-of-the-line notebook computers with the horsepower and features to make traveling executives' dreams come true.

ANTHONY J. LENNON AND JOHN MCDONOUGH

The constant quest for faster portable computers motivates system vendors to deliver desktop system capabilities in portable packages. In this report, we find that, now more than ever, notebook computers offer luxurious features cleverly packed in small cases. We tested 28 high-performance mobile computers with 486DX4 and Pentium processors, a traveling executive's dream. To round out our portables coverage, we also look at Apple's latest high-end Power-Book 540c and AST Research's Ascentia 810N; both represent suitable alternatives to Intel's dominant 486DX4-based systems. We also included the Ascentia 810N to compare its 66-MHz Cyrix 486DX2 processor to notebooks running Intel's 486DX4/75.

Our Roll Call on page 206 describes features found in high-powered desktop machines. Three of these notebooks offer 90-MHz Pentium processors, and most feature 8 MB of RAM, over 500 MB of hard disk space, and 1 MB of video RAM. They all have active-matrix color displays. The Dolch Computer Systems' L-PAC 586 90, the only lunchbox system in this review, offers 1.05 GB of hard disk space and 64 MB of RAM. None of these high-revving systems is what you might call inexpensive: The average as-tested price for a Pentium system in this report is $6500, and the average 486DX4 notebook covered here costs $5000. IBM's ThinkPad 755CD costs as much as an economy car ($7599 as configured for testing), but sales of the ThinkPad line indicate people are willing to pay top dollar for state-of-the-art mobile systems.

How to use this guide

We selected the best high-end notebooks by evaluating speed, screen quality, battery life, price, features, and ease of use.

See the Roll Call on pages 206–209 for features included at this price.

A combination of the weighted scores for performance, screen quality, features, battery life, and usability.

Rates the display's crispness, intensity/color range, and range of viewing angles.

A subjective assessment of the keyboard, pointing device, status lights, and other considerations.

BEST OVERALL AST Ascentia 910N

The Ascentia 910N won our best overall rating mostly due to its excellent battery life and outstanding 10.4-inch color display. This 7-pound system ran for over 6 hours in our Thumper II test. Internally, it can display 256 colors at a resolution of 800 by 600 pixels. The system's 82-key keyboard received above average ratings for its feel and response. Although its overall performance is only average for a 75-MHz 486DX4-based system, its feature set propelled it to the top.

		PRICE AS TESTED	CPU/SPEED	OVERALL SCORE	FEATURES	EASE OF USE	SCREEN QUALITY
BEST	AST Ascentia 910N	$4636	486DX4/75	7.4	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	Digital HiNote CT475	$4699	486DX4/75	7.3	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	HP OmniBook 600CT	$6030	486DX4/100	7.1	▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	NEC Versa M75 High Resolution	$5348	486DX4/75	7.0	▲▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	Twinhead Slimnote 5100T	$4619	486DX4/100	6.9	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	Zenith Z-NoteFlex	$5592	486DX4/75	6.8	▲▲▲	▲▲▲	▲▲▲▲

Behemoths in Disguise

APRIL 1995
BYTE
BEST
HIGH-END
NOTEBOOKS

BEST

LCD
For high-performance notebooks, color is a requirement, and active-matrix screens produce the brightest images. Systems with dual-scan color displays are more affordable, but their displays are not as beautiful.

POINTING DEVICE
It's a matter of opinion, and all pointing devices take some getting used to, but trackballs centered below the keyboard are generally the best for touch-typists. Our testers also liked the eraser-head type of pointing devices, such as IBM's TrackPoints.

REMOVABLE HARD DRIVE
A modular design makes it easy to swap hard drives among notebooks, reconfigure your computer, and upgrade when necessary.

EXTERNAL FLOPPY DRIVE
If you don't need it this trip, leave it at home.

AIR VENTS
Adequate air circulation and heat sinks keep hot-running processors from overheating.

CD-ROM DRIVE
Before you buy a portable, note the CD-ROM drive's speed and placement. Single-speed drives are slow. Often, CD-ROM drives come as external options and might take the place of another component.

PCMCIA SLOTS
All but five notebooks in this report support one Type III card, giving you the option of inserting two Type II cards instead. For ease of setup, make sure your portable supports Card and Socket Services 2.1.

BATTERY
Notebooks with lithium ion batteries generally outlast those with nickel-metal hydride packs, but NiMH batteries are less expensive.

AUDIO
If you need presentation capability, consider a notebook with on-board 16-bit audio, built-in speakers, and a microphone.

BEST OVERALL 486DX4

AST Ascentia 910N

This SL-enhanced 75-MHz system's battery lasts more than 6 hours, nearly 2 hours longer than any other model in this review. For $4635, the 910N ships with 16 MB of RAM and a 510-MB hard drive. Its Hitachi active-matrix display produces beautiful color, saturating the 10.4-inch screen with vibrant images of 256 colors in 800- by 600-pixel resolution. For $4698, the Digital HiNote CT475 offers excellent performance, an excellent screen, and good features and ease-of-use.

PAGE 199

BEST OVERALL PENTIUM

Sceptre Pentium Soundx

The $5995 Pentium Soundx is just the system to meet future demands. Its 90-MHz CPU, 16 MB of RAM, 256-KB direct-mapped write-back secondary cache, 540-MB hard disk, and CTI (Chips & Technologies) local-bus graphics all combine to deliver excellent performance. The Pentium Soundx sports a 10.4-inch active-matrix display, and testers liked the design and response of its 86-key keyboard. Audio features include an on-board stereo FM sound generator, integrated speakers, and a microphone.

PAGE 202

All these systems use low-power 3.3-volt CPUs—either 486DX4 75- or 100-MHz CPUs, P75s, or P90s. Intel's new 75-MHz SL-enhanced Pentium processor adeptly conserves battery life without incurring a major performance hit. And our exclusive battery-life tests indicate that AST's 75-MHz 486DX4-based Ascentia 910N squeezes an impressive 6 hours out of its lithium ion battery. Lithium ion batteries are approximately twice as expensive as NiMH (nickel-metal hydride) batteries.

Our InterMark performance tests employ actual Windows and DOS applications; the BYTE low-level tests stress individual system components, such as graphics and storage subsystems under DOS; and the low-level InterMark suite tests these same subsystems under Windows 3.1. Our Thumper II test determines battery life by simulating a word processing session that might occur on a typical cross-country plane flight. The portables in this report have, as minimum criteria, 75-MHz processors, active-matrix color displays, 8 MB of RAM, 340-MB hard disks, and PCMCIA sockets.

75 MHZ AND FASTER
486DX4 LAB NOTES

T he clock-tripled Intel 486DX4 processors enable systems manufacturers to approach entry-level Pentium performance with minor motherboard upgrades. The 486DX4 processors are available in speeds of up to 100 MHz and feature an integrated 16-KB (level 1) memory cache to reduce overall memory access time. The 486DX4s inherited power-saving features from Intel's 486SL line and they all operate at 3.3 volts, a reduction from the previous 5 volt processors.

In our tests, the average 75-MHz DX4's battery lasted 40 minutes longer than its 100-MHz counterpart (3 hours and 28 minutes versus 2 hours and 47 minutes).

We tested systems that support a second battery pack in place of a standard floppy drive with just one battery. This was particularly detrimental to the Panasonic V41 Multimedia Notebook PC, because its floppy drive replaces its main battery pack. With its NiMH secondary battery, the V41 lasted only 1 hour. To be fair, we then tested it with the main battery, and it ran nearly 2.5 hours. When traveling, though, you might prefer to take out the floppy and go with both battery packs, effectively doubling the battery life.

The 486DX4-based notebooks are well-suited for graphical interfaces. All but the Compaq LTE Elite 4/75CX Model 510 contain 32-bit local-bus graphics.

The NEC Versa M/75 High Resolution and AST's Ascentia 910N support a maximum internal resolution of 800 by 600 pixels, higher than any of the others. But at a maximum internal resolution of 640 by 480 pixels, IBM's ThinkPad 755CD, Panasonic's V41 Multimedia Notebook PC, Hewlett-Packard's OmniBook 4000CT, and Dell's Latitude XP 4100CX can display up to 65,536 colors, whereas the 800- by 600-pixel systems support only 256 colors.

Digital's HiNote CT475 offers the widest viewing angle (30 degrees left, 25 degrees right, and 35 degrees backward) as measured by our HURD (Heads Up Range Device). Developed by NSTL, the HURD tests determine a notebook's maximum viewing angles from left to right and by tilting the screen backward. Compaq's LTE Elite 4/75CX and HP's OmniBook 4000CT provide the most limited viewing angles (about 25 degrees from side to side and 5 degrees to the back).

Integrated pointing devices are standard on most mobile computers. The only notebook without one is Toshiba's T4850CT. Its QuickPort mouse attaches to the right-hand side of the system and is easy to master, but it is of little use to lefties.

The Ascentia 910N and WinBook XP have eraser-head pointing devices, and the Twinhead Slimnote 5100T features an integrated TouchPad, similar to those in Apple's PowerBook 500 series. Many notebook vendors are moving to trackpads; you can also purchase them as accessories (see "Do the Electric Glide" on page 203). Aquiline's Cyclone has an integrated joystick mounted next to the spacebar—great for game addicts.

The best pointing device is a matter of opinion; keyboard quality is also highly subjective. However, we really liked the feel and response of the Austin Business Audio Notebook, T4850CT, and the Slimnote 5100T.

AST ASCENTIA 810N

All the other notebooks in this report run on 75-MHz and higher CPUs. The Ascentia 810N runs on a 66-MHz 486DX2 CPU from Cyrix, but in CPU- and memory-intensive tests like our WordPerfect benchmark, it is a formidable competitor to the others. Configured with 20 MB of RAM and a 510-MB hard disk, the $3948 810N boasts an integrated 8-KB write-back cache, which improves performance by caching memory writes as well as memory reads. The Intel 486DX4 processor's integrated 16-KB cache uses a write-through design that caches memory reads but passes memory writes through to system memory.

The Ascentia 810N is also proficient in the disk-intensive Microsoft FoxPro test. However, it cannot match the 75-MHz 486DX4 notebooks in tests dependent on raw processing power.

Western Digital's RocketChip local bus graphics accelerator enhances the Ascentia 810N's graphics performance, and its NiMH (nickel-metal hydride) battery ran more than 3 hours in our test. You can double that to over 6 hours by replacing the system's floppy disk drive with an optional second battery pack. The Ascentia 810N supports up to 20 MB of RAM and a 510-MB IDE hard drive. It weighs only 6.5 pounds with its battery pack and AC adapter, and supports two Type II or one Type III PC card. The system includes a three-year parts-and-labor warranty plus unlimited technical support.
—Anthony J. Lennon

AST'S 486DX2/66 ASCENTIA 810N VS. 486DX4/75 NOTEBOOKS
The geometric means of three tests' scores comparing the Cyrix 486DX2/66-based AST Ascentia 810N's performance to three of the Intel-based 486DX4/75 notebooks.

	WORDPERFECT 6.0 FOR DOS	LOTUS 1-2-3 R. 3.X FOR DOS	MS FOXPRO 2.5 FOR DOS
AST Ascentia 810N 486DX2/66	9.7	92.6	1526.1
Austin Business Audio Notebook 486DX4/75	9.4	103.8	1538.0
Digital HiNote CT475 486DX4/75	9.8	107.9	1700.5
Toshiba T4850CT 486DX4/75	9.2	104.4	1690.8

To qualify in our Desktop Replacement category, the 486DX4 notebooks had to support a 1024- by 768-pixel resolution externally and at least 32 MB of RAM, a 650-MB hard disk, integrated sound, and an external expansion chassis. Although only eight of these systems could meet these requirements, all could function as a desktop replacement to some degree. On average, the DX4s support up to 32 MB of RAM, more than enough for even the most demanding GUI, and a 785-MB hard disk (removable on all but the T4850CT and TI's TravelMate 4000M). IBM and WinBook include 14.4 Kbps data/9600 bps fax modems, and the feature-rich ThinkPad 755CD also includes a digital answering machine and an integrated infrared port for wireless file transmission.

Many of these portables are equipped with excellent multimedia features. The Panasonic V41, Aquiline's Cyclone, and IBM's ThinkPad all feature internal double-speed CD-ROM drives. You can replace the ThinkPad's CD-ROM drive with an internal floppy drive; the Cyclone features an external floppy drive. All but five systems offer integrated sound chips. The ThinkPad 755CD features a built-in Mwave DSP (digital signal processor) that provides MIDI synthesis and Sound Blaster support. TI's TravelMate 4000M offers an optional double-speed CD-ROM drive.

The Ascentia 910N ascends the throne

BEST OVERALL   AST Ascentia 910N

The Ascentia 910N won our best overall rating mostly due to its excellent battery life and outstanding 10.4-inch color display. This 7-pound system ran for over 6 hours in our Thumper II test. Internally, it can display 256 colors at a resolution of 800 by 600 pixels. The system's 82-key keyboard received above average ratings for its feel and response. Although its overall performance is only average for a 75-MHz 486DX4-based system, its feature set propelled it to the top.

		PRICE AS TESTED	CPU/SPEED	OVERALL SCORE	FEATURES	EASE OF USE	SCREEN QUALITY
BEST	AST Ascentia 910N	$4635	486DX4/75	7.4	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	Digital HiNote CT475	$4698	486DX4/75	7.3	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	HP OmniBook 4000CT	$6030	486DX4/100	7.1	▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	NEC Versa M/75 High Resolution	$5348	486DX4/75	7.0	▲▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	Twinhead Slimnote 5100T	$4619	486DX4/100	6.9	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	Zenith Z-NoteFlex	$5582	486DX4/75	6.8	▲▲▲	▲▲▲	▲▲▲▲

Incredible features for the affluent

DESKTOP REPLACEMENT   IBM ThinkPad 755CD Notebook

The 100-MHz 486DX4-based IBM ThinkPad 755CD Notebook works well as a multimedia desktop replacement system. It boasts a removable 5.25-inch double-speed CD-ROM drive and 16-bit sound. Its 10.4-inch TFT (thin-film transistor) black active-matrix screen displays up to 65,536 colors, and the display's viewing range is excellent. Expansion options include support for up to 40 MB of RAM and an 810-MB removable hard drive. An optional docking station ($710) houses two ISA expansion slots, a 32-bit voice MIDI synthesizer, and a one-third height drive bay. A Western Digital WD90C24 local-bus graphics subsystem speeds graphics performance, and its local-bus hard drive interface made it a high achiever in the disk-intensive benchmark tests.

		PRICE AS TESTED	CPU/SPEED	OVERALL SCORE	FEATURES	EASE OF USE	SCREEN QUALITY
BEST	IBM ThinkPad 755CD Notebook	$7599	486DX4/100	6.8	▲▲▲▲	▲▲▲▲	▲▲▲
RUNNER-UP	NEC Versa M/75 High Resolution	$5348	486DX4/75	6.8	▲▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	WinBook XP	$3799	486DX4/100	6.5	▲▲▲	▲▲▲▲	▲▲▲
RUNNER-UP	Austin Business Audio Notebook	$3319	486DX4/100	6.5	▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	NEC Versa M/75	$4998	486DX4/75	6.3	▲▲▲▲	▲▲▲▲	▲▲▲
RUNNER-UP	Micro International Mint 7500T	$3970	486DX4/100	6.1	▲▲▲	▲▲▲	▲▲▲▲

As the name says, the WinBook's a winner

LOW COST   WinBook XP

With its test configuration price of $3799, the WinBook XP is an excellent value. This 486DX4/100 system ran nearly four and a half hours on its hot-swappable NiMH (nickel-metal hydride) battery. With a traveling weight of 7 pounds, it also features a 14.4 Kbps data/fax modem, integrated 8-bit sound, and a TrackPoint II pointing device. The WinBook's Hitachi screen sports a wide viewing range but only average color quality. An optional docking station ($399) provides 16-bit sound, two ISA expansion slots, and two 3.5-inch drive bays. Maximum RAM is 32 MB, the maximum hard disk size is 700 MB, and its secondary memory cache maxes out at 512 KB.

		PRICE AS TESTED	CPU/SPEED	OVERALL SCORE	FEATURES	EASE OF USE	SCREEN QUALITY
BEST	WinBook XP	$3799	486DX4/100	6.6	▲▲▲	▲▲▲▲	▲▲▲
RUNNER-UP	Austin Business Audio Notebook	$3319	486DX4/100	6.4	▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	Samsung NoteMaster 3945T	$3550	486DX4/75	6.2	▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	Austin Business Audio Notebook	$3239	486DX4/75	6.1	▲▲▲	▲▲▲▲	▲▲▲▲
RUNNER-UP	Micro International Mint 7500T	$3970	486DX4/100	6.1	▲▲▲	▲▲▲	▲▲▲▲
RUNNER-UP	DTK DTN-4T100P	$2600	486DX4/100	5.5	▲▲	▲▲▲	▲▲

How We Tested

PERFORMANCE

We assessed the performance of each system with a suite of BYTE low-level DOS benchmarks and NSTL's applications and low-level InterMark tests for Windows. Prior to testing each system, we installed MS-DOS 6.2 and Microsoft Windows 3.1 onto freshly formatted hard disks.

BYTE's low-level DOS benchmarks measure the performance of specific subsystems: CPU, FPU, memory, graphics, and the hard drive system. NSTL's Windows-based low-level InterMark tests exercise the Windows GDI (Graphical Device Interface), as well as low-level graphics, CPU/FPU, and hard drive performance. The GDI component determines how well a system executes basic graphics calls within Windows. We ran all the Windows-based tests in 640- by 480- by 256-pixel resolution using vendor-supplied graphics drivers.

The applications benchmarks employ popular business applications and portray real-world performance.

FEATURES

We asked each vendor to complete a lengthy questionnaire to give us a detailed description of each system's features and support options. We then weighted each feature and calculated an overall features score (see the Roll Call on page 206).

We focused on three aspects of screen quality: crispness, intensity/color range, and viewing range. We ran tests to gauge the clarity of text in both color and monochrome: horizontal and vertical line placement, the color and gray-scale depths, and the frequency of LCD streaking. We used Sonera Technology's DisplayMate Professional 1.0 to analyze a wide range of display capabilities. To determine color quality, we displayed a color bar on each screen and assigned a score ranging from 1 (worst) to 5 (best). We created a panel of three active-matrix displays judged to be at the low, middle, and high range of quality to assist in our relative judging.

After computing the viewing range using NSTL's HURD, we plotted the group's viewing-angle scores.

We measured battery performance with BYTE's Thumper II system. Thumper mimics real-world use by emulating a typical word processing ses-

sion. Robotic arms and optical sensors detect and control each system's power management scheme. We tested notebooks that support a second battery pack in place of floppy drives with the floppy drive in place.

Prior to testing, we completely drained and recharged each system's battery according to the manufacturer's instructions and then repeated that procedure. We configured each notebook's power management features to spin down the hard disk after 2 minutes of inactivity and shut off the backlighting after 1 minute. We allowed each system to enter standby mode during the test cycle. At intervals, Thumper's robotic arms wake up each system, and each runs until its battery dies. You must interpret the results of any battery-life test, including Thumper II's, with caution; individual usage patterns will affect a system's battery life.

We also assigned a general evaluation score to each system based on our overall subjective impression.

THE CAT'S PAJAMAS: THE POWERBOOK 540c

Apple's latest PowerBook, the 540c, serves up Quadra performance, thanks to its 33-MHz 68LC040 processor. Apple has revamped the design of the PowerBook line to offer a larger (10-inch diagonal) active-matrix color display, a built-in Ethernet port, and a solid-state trackpad that replaces the trackball; however, the PowerBooks still fulfill Apple's all-in-one design goal. At 640- by 400-pixel resolution, the 540c supports 32,768 colors; at 640 by 480, it supports 256 colors. The power key is now on the keyboard, where it belongs, instead of at the back of the computer. Shutting the lid automatically switches the notebook into sleep mode.

A duplicate accessory bay can hold a second lithium battery or a PCMCIA card cage. A two-battery setup effectively doubles battery life to 5 hours and 34 minutes, yet the 540c weighs only 7 pounds. The PCMCIA card cage can hold up to two Type II PC cards or a single Type III card, and you can hot dock PC cards while the 540c is awake. The latest version of PC Exchange (Apple's system software for mounting DOS volumes) now has the capability to mount and access DOS-formatted PCMCIA storage devices. Such devices can be either flash memory or a hard drive. It may be a tough choice whether to go for the extra battery life or the accessibility to PCMCIA devices. A major plus of the 540c is its easily removable CPU card. You'll be able to transform the 68LC040/33 540c into a PowerPC as soon as that option becomes available. The 540c as tested includes 12 MB of RAM and a 500-MB hard drive. It lists for $4999.

USABILITY

We also assessed the quality of each keyboard, concentrating specifically on key placement. We worked extensively with the notebooks to see how comfortable they became after extended use, and we rated the response and feel of the keys. We also evaluated pointing devices for both right- and left-handed people, and considered the usefulness of the LCD status indicators. We pored over the documentation, taking into account both novice and expert needs.

Contributors

Anthony J. Lennon, *Project Manager/NSTL*
John McDonough, *Technical Writer/NSTL*
Maggi Bender, *Tester/NSTL*

The Lab Report is an ongoing collaborative project between BYTE magazine and National Software Testing Laboratories (NSTL). BYTE magazine and NSTL are both operating units of McGraw-Hill, Inc. Contact NSTL on the Internet at editors@nstl.com; at NSTL, Inc., Plymouth Corporate Center, Plymouth Meeting, PA 19462; or at (610) 941-9600. Contact BYTE via the Internet or BIX at editors@bix.com or at (603) 924-2643.

ZEOS Meridian™ Notebooks—
Hot Performance At A Light Price.

Pentium Lab Notes

The emergence of 3.3-volt, SL-enhanced Pentium processors makes Pentium-based notebooks very nice traveling companions. With the exception of Dolch's L-PAC 586 90 MHz, which doesn't support battery power at all, all these Pentium notebooks provide NiMH battery packs, and all (including the 486DX4 systems) contain 3.3-volt CPUs. AT&T's Globalyst 250 and Toshiba's T4900CT use 75-MHz CPUs and both lasted approximately 3 hours in our battery-life test (the Globalyst ran 2 hours and 42 minutes, and the T4900CT persisted half an hour beyond that). Sceptre's 90-MHz notebook also ran for nearly 3 hours (2 hours, 54 minutes). With a 256-KB direct-mapped write-through secondary memory cache, the Globalyst outperforms Toshiba's T4900CT in many memory-intensive applications, such as the Lotus 1-2-3 mathematical suite. Due to its local-bus graphics subsystem, the Globalyst holds an ace of spades in any graphics-intensive test.

With 90-MHz processors, 256-KB secondary caches, and local-bus graphics, the Sceptre Pentium Soundx and the Mitsuba Ninja-P90 excel in overall performance. Dolch's L-PAC has an inferior graphics design, which hurt its overall performance score. The L-PAC's 16-bit graphics circuitry lies on the same multifunction I/O card as its PCMCIA controller, and the notebook's graphics performance pales considerably when compared to any 32-bit local bus sub-

Clockwise from top left: Mitsuba's Ninja-P90, Sceptre's Pentium Soundx, AT&T's Globalyst 250, and Toshiba's T4900CT. The Ninja-P90 and the Pentium Soundx look nearly identical.

system's. However, the L-PAC does feature a 32-bit PCI (Peripheral Component Interconnect) hard drive, making it a high-achiever in the disk-intensive FoxPro benchmarks. The Globalyst 250 performed poorly in the DOS-based FoxPro test. We concluded the Globalyst must have gone into local standby mode, even though we set the power management to Disabled. Toshiba's T4900CT features a local-bus IDE hard drive and outperforms the Globalyst in our low-level file read/write tests, as well as in

the WordPerfect for Windows file I/O benchmark.

All the Pentiums have active matrix color LCDs capable of displaying up to 256 colors at a resolution of 640 by 480 pixels. We awarded the best color-quality ratings to Mitsuba's Ninja and to Sceptre's Soundx. The Globalyst 250 produced above average colors and offered the widest viewing angle. The Dolch L-PAC also offers a wide 30-degree viewing range from side to side, but you can't tilt its screen backwards.

BEST OVERALL — SCEPTRE PENTIUM SOUNDX

With its 90-MHz processor, the Sceptre Pentium Soundx delivers state-of-the-art performance. The notebook's large 10.4-inch Toshiba display produces vibrant, fully saturated colors. Battery life is excellent for a Pentium system: With its NiMH (nickel-metal hydride) battery pack, it ran almost three hours in our Thumper II battery test. Our testers enjoyed the feel and response of its 86-key keyboard. The system's integrated 16-mm trackball, though conveniently placed in the middle of the wrist rest, is difficult to control. Audio features include an on-board 16-bit stereo FM sound generator, two integrated speakers, and a microphone. It supports up to 40 MB of RAM, a 720-MB hard drive, and an optional expansion chassis ($595) with two 16-bit slots and two 3.5-inch drive bays. An included remote-control mouse (shown at right) gives you mouse control within a 45-foot range. Other portables support third-party remote-control devices.

	PRICE AS TESTED	CPU	OVERALL SCORE	FEATURES	EASE OF USE	SCREEN QUALITY
Sceptre Pentium Soundx	$5995	P90	8.0	▲▲▲	▲▲▲▲	▲▲▲▲
Mitsuba Ninja-P90	$4995	P90	7.7	▲▲▲	▲▲▲▲	▲▲▲▲
AT&T Globalyst 250	$5535	P75	7.6	▲▲▲▲	▲▲▲▲	▲▲▲▲
Dolch L-PAC 586 90 MHz	$8975	P90	7.2	▲▲	▲▲▲▲	▲▲▲
Toshiba T4900CT	$7499	P75	6.8	▲▲▲	▲▲▲	▲▲▲

Key: Ratings from 1 to 4: ▲ is the lowest, ▲▲▲▲ is the highest.

75- to 90-MHz Portables

Dolch

Dolch's L-PAC 586 90 MHz is a unique system, designed primarily for engineers. Its 16-pound lunch-box-style design (10.9 by 15.7 by 6.3 inches) is reminiscent of early luggables, but it is intended to be a true desktop replacement rather than a mobile companion. You can easily install up to three 16-bit ISA adapters and enjoy a full-size, 101-key detachable keyboard with a dedicated numeric keypad. It has fold-down feet on the bottom and a handle on top, but you won't want to carry the L-PAC too far. It doesn't run on battery power and lacks an integrated trackball and built-in 16-bit sound; however, it does provide two conveniently located Type II PCMCIA slots, and you can configure it with up to 64 MB of RAM and a 1.05-GB hard drive. The Pentium-

powered L-PAC holds an FCC Class A rating and a steep price tag. Its $8975 review-configuration price includes 32 MB of RAM and a 1.05 GB IDE hard drive.

By The Time You Read This....

Between the time we began testing and the time you read this, vendors will have upgraded their lines with yet more impressive features. Two notebooks in particular caught our attention: Tadpole Technology's Pentium-powered P1000 (it won our "Best Portable" award last November at Comdex) and Sharp Electronics' PC-8900 multimedia notebook (particularly impressive due to its advanced LCD technology).

Tadpole designs workstations, and the P1000 is no exception. It runs on Intel's 3.3-volt 100-MHz 815/100 CPU, has an integrated 256-KB level 2 write-back cache, a PCI bus, PCMCIA and SCSI interfaces, and maxes out at 128 MB of DRAM. The P1000 has an active-matrix 10.4-inch TFT (thin-film transistor) display panel, internal microphone, speakers, and audio in/out ports. With a battery, the P1000 weighs 7.5 pounds and has a base MSRP (with 8 MB of RAM and a 340-MB hard drive) of $7495.

Sharp's PC-8900 runs on an Intel 486DX4/75 processor and has an infrared port, a Type III PCMCIA slot, and an Alps Electric GlidePoint (see "Do the Electric Glide"). Its 10.4-inch active-matrix screen consumes 20 to 30 percent less power than Sharp's previous LCDs, runs off a VESA local bus, and internally supports 640 by 480 pixels by 256 colors. Its multimedia features include 16-bit sound; integrated speakers and microphone; and the ability to mix, edit, and synthesize audio files. Weighing in at 6.2 pounds, its target price is $4999.

—Selinda Chiquoine

DO THE ELECTRIC GLIDE

The GlidePoint from Alps Electric offers a unique alternative to traditional point-and-click devices. By gliding your fingertip across the GlidePoint's 2.52- by 1.9-inch nonmoving surface, you control your screen's cursor. Because it has no moving parts, the unit never needs cleaning. The wafer-thin pad presents a major threat to standard pointing devices; it's rapidly gaining marketshare.

Alps Electric applies field-distortion sensing technology, a form of capacitance-sensing licensed from Cirque (Salt Lake City, UT). Two layers of fine electrical conductors in a planar grid create an electrical field on the GlidePoint's sealed surface. When your finger (pencils and erasers don't work) touches the surface, you distort the pad's electrical field, allowing it to track your fingertip as it moves across the grid. You can "click" by tapping your finger lightly on the pad, but the GlidePoint has three buttons as well. The driver Alps provides is simple to install, and with it you can program the third button. The Glide-Point also works with Microsoft mouse drivers.

Although Alps touts it as a "natural finger pointing move-

ment," expect to spend a little time getting adjusted to the new technology. If you're not unhappy with your current trackball or mouse, the GlidePoint might not be worth the $96 investment. The Twin-head Slimnote 5100T and the Apple Power-

The Alps Electric GlidePoint pointing device

Book 500 series currently incorporate trackpads; expect to see more notebook vendors adopting the trackpad technology soon.

Epson America, Inc., has announced that all the systems in its ActionNote 800 series will offer internal TouchPads, and WinBook Computer's WinBook XP line will offer internal Touch-Pads or trackballs (centered in front of the spacebar) as a pre-installed option.   **—John McDonough**

HONORABLE MENTIONS

The Panasonic V41 Multimedia Notebook
PC features an internal double-speed CD-ROM drive. The drive lies beneath the keyboard, which opens with a conveniently placed latch on the notebook's front. An adjacent slot stores a CD. The **IBM ThinkPad 755CD Notebook** and the **Aquiline Cyclone** also feature internal double-speed CD-ROM drives. However, you must remove IBM's floppy drive to install the CD-ROM drive, and the Aquiline ships with an external—rather than internal—floppy drive.

The Digital Equipment HiNote CT475's LCD status display is highly visible. It is located to the right of the screen, unlike the majority of systems which place their status indicators on the display hinge or above the keyboard. Status indicators are especially useful if they display a multilevel battery gauge, and the HiNote's does.

The Compaq LTE Elite 4/75CX Model 510 features a built-in AC adapter, eliminating a bulky accessory. The **IBM ThinkPad 755CD Notebook** includes a 12-foot power cord and a 3-prong plug that pops out of the AC adapter's box, so there's no need for a power cord between the adapter and the wall socket. However, be careful: IBM discovered that the adapters labeled model AA19210, with a date code of 9452 or lower, may present the risk of electrical shock. Although no shocks have been reported, IBM is voluntarily recalling these units.

Two innovative power designs. The ThinkPad's is on the left and the LTE Elite's is on the right.

Dubious Achievements

◄ **The Compaq LTE Elite 4/75CX Model 510's** integrated trackball is awkwardly placed. You must take your right hand off the keyboard and grasp the right-hand side of the display to use the trackball. However, its molded buttons (shown below, left) located on the back of the display panel conveniently fit your fingers.

The Zenith Z-NoteFlex received an above-average rating for the feel and response of its 82-key keyboard. However, we found the symbols on the Home, End, PageUp, and PageDown keys to be very cryptic and of little use to novices. And because there is no hard-copy User's Guide, you must consult the system's Windows-based on-line documentation to decipher the function keys' function.

►

◄ **The placement of the trackball** on the AT&T Globalyst 250, NEC Versa M/75, and NEC Versa M/75 High Resolution (whose exteriors look identical except for their color and displays) makes it difficult to use. These trackballs are located on the front of the system in between buttons for selecting and dragging.

SCEPTRE ®

THE PEOPLE WHO BROUGHT YOU THE AWARD WINNING MONITORS, HAVE A TERRIFIC NEW SERIES OF PENTIUM AND 486 NOTEBOOK COMPUTERS!

SOUNDX™ Multimedia
Audio Notebook Computer

3000 Series
▼ Pentium to 90Mhz
▼ Upgradeable to 40MB RAM
▼ Modular Design
▼ 10.4" TFT/ DSTN
▼ PCI Bus
▼ 16-bit Stereo (SoundBlaster Pro)
▼ Docking Station Port
▼ PCMCIA Interface
▼ INFRARED Remote Link

1000 Series
▼ 486 up to DX4/100Mhz
▼ Modular Design
▼ 10.3" Dual Scan Color
▼ 16-bit Stereo (SoundBlaster Pro)
▼ Docking Station Port
▼ PCMCIA Interface

intel inside

PCMCIA Interface

I/O Ports

Battery Status Indicator

CPU/Memory/HDD
Easy to upgrade & service

Docking Station

The **Soundx™ Multimedia Notebook Computer** is an affordable, lightweight, versatile solution to all mobile high performance needs. The unique modular design allows for upgradeability of the CPU, memory, hard drive and LCD screen. Other valuable features include larger viewing screen, PCMCIA interface, energy saving power management, built-in trackball, system status LCD indicator, and 16-bit sound chip. **Soundx™** represents **Sceptre's** commitment to advanced technology through environmentally safe products.

SCEPTRE
TECHNOLOGIES
16800 E. Gale Ave. City of Industry, CA 91745 USA
Tel:(818) 369-3698   Fax:(818) 369-3488
Malaysia: 03-703-3133   Canada: 604-244-7798
SM-1P5-A

Intel Inside logos are trademarks of Intel Corporation. All trademarks are registered with their respective companies.

Circle 120 on Inquiry Card (RESELLERS: 121).

ROLL CALL OF NOTEBOOKS TESTED

VENDOR	MODEL NAME	PERFORMANCE RATING	BATTERY LIFE (HR:MIN:SEC) AS TESTED	FEATURES/ EASE OF USE	SCREEN QUALITY	PRICE AS TESTED (US$ MSRP)	CPU/MHZ	RAM AS TESTED/ MAX. 32-BIT RAM (MB)	BIOS VENDOR
Aquiline, Inc.	Cyclone[2]	▲▲	2:12:26	▲▲/▲▲▲	▲▲	$3795	486DX4/100	8/64	AMI
AST Research, Inc.	Ascentia 910N	▲▲	6:09:07	▲▲▲/▲▲▲	▲▲▲▲	$4635	486DX4/75	16/32	AST
Austin Direct	Business Audio Notebook/75	▲▲	2:22:06	▲▲▲/▲▲▲▲	▲▲▲▲	$3239	486DX4/75	8/32	Phoenix
Austin Direct	Business Audio Notebook/100	▲▲▲	2:03:56	▲▲▲/▲▲▲▲	▲▲▲▲	$3319	486DX4/100	8/32	Phoenix
Compaq Computer Corp.	LTE Elite 4/75CX Model 510[3]	▲▲	2:55:15	▲▲▲▲/▲▲▲	▲▲▲	$5799	486DX4/75	8/24	Compaq
Dell Computer Corp.	Latitude XP 4100CX	▲▲▲	3:57:50	▲▲▲/▲▲▲▲	▲▲▲	$4499	486DX4/100	8/36	Phoenix/Dell
DFI	MediaBook 5110T	▲▲▲	1:39:15	▲▲▲/▲▲▲	▲▲▲	$4740	486DX4/100	36/36	Award
Digital Equipment Corp.	HiNote CT475	▲▲	4:09:10	▲▲▲/▲▲▲	▲▲▲▲	$4698	486DX4/75	8/20	System Soft
DTK Computer, Inc.	DTN-4T100P[2]	▲▲	2:21:00	▲▲/▲▲▲	▲▲	$2600	486DX4/100	8/36	AMI
Hewlett-Packard Co.	OmniBook 4000CT	▲▲▲	3:24:26	▲▲▲/▲▲▲▲	▲▲▲▲	$6030	486DX4/100	8/32	Phoenix
IBM Personal Computer Co.	ThinkPad 755CD Notebook	▲▲▲	3:19:01	▲▲▲▲/▲▲▲▲	▲▲▲	$7599	486DX4/100	8/40	IBM
Micro International, Inc.	Mint 7500T	▲▲▲	2:07:29	▲▲▲/▲▲▲	▲▲▲▲	$3970	486DX4/100	8/36	AMI
NEC Technologies, Inc.	Versa M/75	▲▲▲	2:59:07	▲▲▲▲/▲▲▲▲	▲▲▲	$4998	486DX4/75	16/40	NEC/Phoenix
NEC Technologies, Inc.	Versa M/75 High Resolution	▲▲▲	3:37:37	▲▲▲▲/▲▲▲▲	▲▲▲▲	$5348	486DX4/75	16/40	NEC/Phoenix
Panasonic Personal Computer Co.	V41 Multimedia Notebook PC	▲▲▲	1:00:05[4]	▲▲▲/▲▲	▲▲▲	$6599	486DX4/100	8/32	IBM/Phoenix
Samsung Electronics America	NoteMaster 3945T[2]	▲▲▲	2:38:53	▲▲/▲▲▲▲	▲▲▲▲	$3550	486DX4/75	8/20	Phoenix
Texas Instruments, Inc.	TravelMate 4000M DX4/100 Color[2]	▲▲▲	3:02:33	▲▲▲/▲▲▲▲	▲▲▲	$5499	486DX4/100	8/20	Phoenix
Toshiba America Information Systems	T4850CT	▲▲	3:23:50	▲▲▲/▲▲▲▲	▲▲▲	$5349	486DX4/75	8/24	Toshiba
Twinhead Corp.	Slimnote 5100T	▲▲▲	4:11:30	▲▲▲/▲▲▲	▲▲▲▲	$4619	486DX4/100	16/32	Phoenix
Unisys Corp.	PW2 Travel Asset	▲▲	2:36:30	▲▲▲/▲▲▲	▲	$4730	486DX4/75	8/20	Phoenix
WinBook Computer Corp.	WinBook XP	▲▲▲	4:21:44	▲▲▲/▲▲▲▲	▲▲▲	$3799	486DX4/100	8/32	Phoenix
Zenith Data Systems	Z-NoteFlex	▲▲	3:55:58	▲▲▲/▲▲▲	▲▲▲▲	$5582	486DX4/75	8/24	ZDS/Phoenix
Zeos International, Ltd.	Meridian 800	▲▲▲	2:37:41	▲▲/▲▲▲	▲▲▲▲	$3995	486DX4/100	8/20	Phoenix
AT&T Global Information Solutions	Globalyst 250	▲▲▲▲	2:42:14	▲▲▲▲/▲▲▲▲	▲▲▲▲	$5535	Pentium/75	8/40	Phoenix
Dolch Computer Systems	L-PAC 586 90 MHz[1,3]	▲▲▲▲	N/A[6]	▲▲/▲▲▲▲	▲▲▲	$8975	Pentium/90	32/64	AMI
Mitsuba Corp.	NINJA-P90[1]	▲▲▲▲	2:16:53	▲▲▲/▲▲▲▲	▲▲▲▲	$4995	Pentium/90	8/36	Award
Sceptre Technologies, Inc.	Pentium Soundx[1]	▲▲▲▲	2:54:00	▲▲▲/▲▲▲▲	▲▲▲▲	$5995	Pentium/90	16/40	Phoenix
Toshiba America Information Systems	T4900CT	▲▲▲▲	3:15:10	▲▲▲/▲▲▲	▲▲▲	$7499	Pentium/75	8/40	Toshiba

= BYTE Best.   [1] Not SL enhanced.   [2] Not flash ROM upgradable   [3] No VESA local bus video

[4] Panasonic V41 Multimedia Notebook PC with main battery pack - 2:26:46   [5] Now shipping with 10.4-inch screen.   [6] Comes with built-in slots.

MAX. DISPLAY RESOLUTION/ TOTAL COLORS AT MAX. RESOLUTION	DIAGONAL SCREEN SIZE (INCHES)	SCREEN MFR.	GRAPHICS ACCELERATOR MFR. AND MODEL	GRAPHICS MEMORY AS TESTED/ BUS WIDTH (BITS)	SIMULTANEOUS INTERNAL & EXTERNAL DISPLAY/KEY-ACTIVATED TOGGLE FOR DISPLAY SELECTION	PORTABILITY HEIGHT × WIDTH × DEPTH (INCHES)	TRAVELLING WEIGHT (LBS.)	BATTERY TYPE	ESTIMATED BATTERY LIFE (HRS.)	BATTERY SUPPORTS MULTIPLE BATTERY PACKS/HOT SWAPPING	RECHARGES WHILE SYSTEM IS ON
640×480/256	9.5	Sanyo	Chips & Technologies 65522	1 MB VRAM/32	Yes/No	2.8 × 11.8 × 8.3	8.5	NiMH	3.5	No/No	Yes
800×600/256	10.4	Hitachi	Western Digital WD90C24A	1 MB VRAM/32	Yes/Yes	1.8 × 11.5 × 8.5	7.0	Li-Ion	6–10	No/Yes	Yes
640×480/256	9.5	NEC	Western Digital WD90C24	1 MB DRAM/16	Yes/Yes	2 × 11.1 × 8.6	8.0	NiMH	3	No/No	Yes
640×480/256	9.5	NEC	Western Digital WD90C24	1 MB DRAM/16	Yes/Yes	2 × 11.1 × 8.6	8.0	NiMH	3	No/No	Yes
640×480/256	9.5	Compaq	Compaq (custom)	1 MB DRAM/32	Yes/Yes	2 × 11.8 × 8.9	7.0	NiMH	4.5	No/Yes	Yes
640×480/64K	9.5	INP	Western Digital WD90C24A2	1 MB VRAM/32	Yes/Yes	2 × 11 × 8.7	7.0	Li-Ion	6–8	No/Yes	Yes
640×480/256	9.5	NEC	Cirrus Logic CL-GD6440	1000 DRAM/32	Yes/Yes	2.2 × 11.2 × 8.8	8.5	NiMH	2	Yes/Yes	No
640×480/256	9.5	Toshiba	Chips & Technologies 65540	1 MB VRAM/32	Yes/Yes	1.7 × 11 × 8.5	7.0	NiMH	3–5	No/Yes	Yes
640×480/256	10.3	Sanyo	Cirrus Logic CL-GD6440	1 MB VRAM/32	Yes/Yes	2 × 11.1 × 9.5	8.0	NiMH	2–3	No/No	Yes
640×480/64K	10.4	Hitachi	Western Digital WD90C24	1 MB VRAM/32	Yes/Yes	1.9 × 11.6 × 8.9	8.0	NiMH	3	Yes/No	Yes
640×480/64K	10.4	DTI	Western Digital WD90C24	1 MB DRAM/32	Yes/Yes	2.2 × 11.7 × 8.3	7.5	NiMH	3–9	No/Yes	Yes
640×480/256	9.5	NEC	Cirrus Logic CL-GD6440	1 MB DRAM/32	Yes/Yes	2 × 11 × 9.4	8.5	NiMH	3	No/No	No
640×480/256	9.5	NEC	Chips & Technologies 65545	1 MB DRAM/32	Yes/Yes	2.1 × 11.7 × 9.5	8.0	NiMH	2.5	Yes/Yes	Yes
800×600/256	9.5	NEC	Chips & Technologies 65545	1 MB DRAM/32	Yes/Yes	2.1 × 11.7 × 9.5	8.0	NiMH	2.5	Yes/Yes	Yes
640×480/64K	10.4	Matushita	Chips & Technologies 65545	1000 DRAM/32	Yes/No	2.3 × 11.7 × 9.3	10.0	NiMH	4.5	Yes/Yes	Yes
640×480/256	9.4	Samsung	Western Digital WD90C24A	1 MB VRAM/32	Yes/Yes	2.2 × 11.3 × 8.9	7.0	NiMH	3	No/Yes	Yes
640×480/256	9.5	Samsung	Cirrus Logic CL-GD6440	1 MB DRAM/32	Yes/Yes	2.1 × 11 × 8.5	7.0	NiMH	3–4	No/No	Yes
640×480/256	10.4	INP	Western Digital WD90C24A	1 MB DRAM/32	Yes/Yes	2.2 × 11.7 × 8.3	7.5	NiMH	2	No/No	Yes
640×480/256	9.4	Sharp	Western Digital WD90C24A2	1 MB DRAM/32	Yes/Yes	1.6 × 11.2 × 8.6	8.0	NiMH	3	Yes/Yes	No
640×480/256	9.5	NEC	Western Digital WD90C24	1000 DRAM/32	Yes/Yes	1.3 × 11 × 8.5	7.5	NiMH	1.5	No/Yes	Yes
640×480/256	9.4	Hitachi	Western Digital WD90C24C	1 MB DRAM/32	Yes/Yes	1.7 × 11.3 × 8.5	7.0	NiMH	2.5	No/Yes	Yes
640×480/256	9.5	NEC	Western Digital WD90C24A	1 MB VRAM/32	Yes/Yes	2.2 × 11.9 × 8.7	8.0	NiMH	2–4	Yes/Yes	Yes
640×480/256	9.5	Toshiba	Cirrus Logic CL-GD6440	1 MB DRAM/32	Yes/No	1.9 × 11.7 × 8.9	7.3	NiMH	3.5	No/No	No
640×480/256	9.5[5]	NEC	Chips & Technologies 65545	1 MB VRAM/32	Yes/Yes	2.1 × 11.7 × 9.3	8.5	NiMH	3–5	Yes/Yes	Yes
640×480/256	9.5	NEC	Chips & Technologies 65535	512 VRAM/16	Yes/No	10.9 × 15.7 × 6.25	16.0	N/A[6]	N/A[6]	N/A[6]	N/A[6]
640×480/256	10.3	Sharp	Chips & Technologies 65545	1 MB DRAM/32	Yes/Yes	2 × 11 × 8.5	7.5	NiMH	1.5	No/No	Yes
640×480/256	10.4	Toshiba	Chips & Technologies 65545	1 MB VRAM/32	Yes/Yes	2 × 11 × 8.6	7.5	NiMH	2	No/No	Yes
640×480/256	10.4	INP	Western Digital WD90C24A	1 MB DRAM/32	Yes/Yes	2.2 × 11.7 × 8.3	7.5	NiMH	2	No/No	Yes

N/A = Not applicable   INP = Information not provided   **Key:** Ratings are from 1 to 4: ▲ is the lowest, ▲▲▲▲ is the highest.

Note: All systems support 220/240 autoswitching and autodetecting AC voltage.

ROLL CALL OF NOTEBOOKS TESTED

VENDOR	MODEL NAME	HARD DRIVE MODEL	HARD DRIVE INTERFACE TYPE/ AVERAGE ACCESS TIME (MS)	HARD DRIVE CONTROLLER CACHE (KB)	HARD DRIVE FORMATTED CAPACITY (MB)/LARGEST AVAILABLE (MB)	REMOVABLE HARD DRIVE/ CD-ROM	STANDARD 3.5-INCH FLOPPY DRIVE	PARALLEL PORT	NUMBER OF PCMCIA SLOTS TYPE II/ TYPE III
Aquiline, Inc.	Cyclone[2]	Toshiba MK246FC	VESA/11	256	540/1 GB	Yes/Yes (1x)	External	EPP	2/1
AST Research, Inc.	Ascentia 910N	Quantum DA 514AT	IDE/19	16	510/720	Yes/Opt.	Internal	EPP	2/1
Austin Direct	Business Audio Notebook/75	Toshiba MK2428FC	VESA/9.7	256	524/810	Yes/None	Internal	EPP	2/1
Austin Direct	Business Audio Notebook/100	Toshiba MK2428FC	VESA/9.7	256	524/810	Yes/None	Internal	EPP	2/1
Compaq Computer Corp.	LTE Elite 4/75CX Model 510[3]	Compaq (Custom)	IDE/17	N/A	514/510	Yes/Opt.	Internal	EPP	2/1
Dell Computer Corp.	Latitude XP 4100CX	Seagate ST9655AG	VESA/16	128	524/810	Yes/Opt.	Internal	EPP	2/1
DFI	MediaBook 5110T	Toshiba MK2526FC	IDE/12	128	503/503	Yes/Opt.	Internal	EPP	1/1
Digital Equipment Corp.	HiNote CT475	Toshiba MK2326FCV	VESA/13	128	340/520	Yes/None	Internal	EPP	2/1
DTK Computer, Inc.	DTN-4T100P[2]	IBM DSAA-2540	VESA/INP	96	520/520	Yes/Opt.	Internal	1-way	2/1
Hewlett-Packard Co.	OmniBook 4000CT	Toshiba MK2326FCH	IDE/12	128	340/520	Yes/None	Removable Internal	EPP	2/1
IBM Personal Computer Co.	ThinkPad 755CD Notebook	IBM DHAA-2540	VESA/14	64	520/810	Yes/Yes (1x)	Removable Internal	EPP	2/1
Micro International, Inc.	Mint 7500T	IBM OEM 520	VESA/13	N/A	520/820	Yes/Opt.	Internal	1-way	2/1
NEC Technologies, Inc.	Versa M/75	IBM DHAA-2540	IDE/15	64	540/810	Yes/Opt.	Removable Internal	EPP	2/1
NEC Technologies, Inc.	Versa M/75 High Resolution	IBM DHAA-2540	IDE/15	64	540/810	Yes/Opt.	Removable Internal	EPP	2/1
Panasonic Personal Computer Co.	V41 Multimedia Notebook PC	Seagate ST9550AG	IDE/16	120	450/680	Yes/Yes (1x)	Removable Internal	EPP	2/1
Samsung Electronics America	NoteMaster 3945T[2]	Toshiba MK2326FCH	IDE/12	128	340/500	Yes/None	Internal	2-way	1/1
Texas Instruments, Inc.	TravelMate 4000M DX4/100 Color[2]	Seagate ST9655AG	IDE/16	128	524/524	No/Yes (2x)	Internal	EPP	2/1
Toshiba America Information Systems	T4850CT	Toshiba MK2428FC	VESA/12	512	500/772	No/None	Internal	EPP	1/1
Twinhead Corp.	Slimnote 5100T	Hitachi DK211A-51	VESA/12	256	500/500	Yes/Opt.	Removable Internal	EPP	2/1
Unisys Corp.	PW2 Travel Asset	IBM DHAA-2405	IDE/19	256	340/340	Yes/Opt.	External	EPP	2/1
WinBook Computer Corp.	WinBook XP	Seagate ST9655AG	VESA/16	512	520/700	Yes/Opt.	Internal	EPP	2/1
Zenith Data Systems	Z-NoteFlex	Seagate ST9655AG	VESA/<16	16	520/700	Yes/Opt.	Removable Internal	EPP	2/2
Zeos International, Ltd.	Meridian 800	Areal A340	IDE/15	32	350/350	Yes/None	Internal	EPP	2/1
AT&T Global Information Solutions	Globalyst 250	IBM DHAA-2540	IDE/13	32	520/810	Yes/Opt.	Removable Internal	EPP	2/1
Dolch Computer Systems	L-PAC 586 90 MHz[1,3]	Conner CFA1080A	PCI/12.5	256	1050/1050	No/Opt.	Internal	EPP	2/1
Mitsuba Corp.	NINJA-P90[1]	IBM N540I	PCI-IDE/12	N/A	520/810	No/Opt.	Internal	2-way	2/1
Sceptre Technologies, Inc.	Pentium Soundx[1]	IBM DBOA-2540	IDE/12	256	540/720	No/Opt.	Internal	EPP	2/1
Toshiba America Information Systems	T4900CT	IBM DVAA-2810	VESA/15	512	772/772	No/None	Internal	EPP	1/1

= BYTE Best. [1] Not SL enhanced. [2] Not flash ROM upgradable [3] No VESA local bus video

[7] P = parts, L = labor, C = shipping to vendor, R = return shipment

BUILT-IN 16-BIT SOUND	BUILT-IN MICROPHONE/ EXTERNAL MICROPHONE PORT	EXTERNAL SPEAKER PORT/SPEAKER VOLUME CONTROL	PASSWORDS POWER-ON/ KEYBOARD/SETUP UTILITY	NUMBER OF KEYS/KEY TRAVEL (MM)	EXPANSION CHASSIS/ DOCKING STATION AND PRICES, IF OPTIONAL	WARRANTY LENGTH (YRS) & COVERAGE[7]	TOLL-FREE PHONE NUMBER	PHONE NUMBER	INQUIRY NUMBER
Yes	Yes/Yes	Yes/Yes	Yes/No/No	81/5	No/No	5 PLR	(800) 370-3322	(518) 272-0421	1346
No	No/No	No/No	Yes/Yes/Yes	82–83/3	Yes: $149/Yes: $489	3 PLCR	(800) 876-4278	(714) 727-4141	1347
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	85/4	No/Yes: $259	1 PLR	(800) 752-1577	(512) 339-3500	1348
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	85/4	No/Yes: $259	1 PLR	(800) 752-1577	(512) 339-3500	1349
No	No/No	No/Yes	Yes/Yes/Yes	82/2	Yes: $329–$399/Yes: $699	3 PLCR	(800) 345-1518	Call Local Compaq Dealer	1350
No	No/No	No/No	Yes/No/Yes	85/3	Yes: $199/No	3 PL	(800) 613-3355	(512) 338-4400	1351
No	Yes/Yes	Yes/Yes	Yes/No/Yes	86/3	Yes: $100/Yes: $495	1 PLCR	(800) 808-4334	(916) 568-1234	1352
No	No/No	No/No	Yes/No/Yes	82/3	Yes: $149/No	3 PLCR	(800) 722-9332	(603) 884-5111	1353
No	Yes/Yes	Yes/Yes	Yes/No/Yes	86/INP	No/Yes: $600	1 PL	(800) 289-2385	(818) 810-0098	1354
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	85/2.5	Yes: $199/No	3 PLCR	(800) 443-1254	(503) 715-2004	1355
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	85/3	Yes: $109–$299/ Yes: $399–$710	3 PLCR	(800) 772-2227	(914) 766-1900	1356
No	Yes/Yes	Yes/Yes	Yes/No/Yes	86/4	No/Yes: $490	1 PLR	(800) 967-5667	(713) 495-9096	1357
Yes	Yes/Yes	Yes/Yes	Yes/Yes/No	83/3	Yes: $199/Yes: $629	3 PL	(800) 632-4636	(408) 433-1200	1358
Yes	Yes/Yes	Yes/Yes	Yes/Yes/No	83/3	Yes: $199/Yes: $629	3 PL	(800) 632-4636	(408) 433-1200	1359
Yes	No/Yes	Yes/Yes	Yes/Yes/No	84/3	Yes: NPA/Yes: $649	1 PLR	(800) 742-8086	(201) 271-3182	1360
Yes	Yes/Yes	Yes/No	Yes/No/Yes	84/3	No/No	2 PLC	(800) 726-7864	(201) 229-4000	1361
Yes	Yes/Yes	Yes/Yes	Yes/No/No	84/3	No/Yes: $849	1 PLR	(800) 848-3927	(817) 771-5856	1362
Yes	Yes/Yes	Yes/Yes	Yes/Yes/No	82/3	No/Yes: $649	3 PLR	(800) 334-3445	(714) 583-3000	1363
Yes	Yes/Yes	Yes/Yes	Yes/Yes/Yes	84/3	Yes: $199/Yes: $599	1 PLCR	(800) 995-8946	(408) 945-0808	1364
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	85/2.5	Yes: $190/Yes: $500	3 PL	(800) 448-1424	(408) 434-2848	1365
No	Yes/No	Yes/Yes	Yes/Yes/Yes	82/3	No/Yes: $399	1 PLR	(800) 468-2162	N/A	1366
Yes	Yes/Yes	Yes/Yes	Yes/No/Yes	82/2	Yes: $209/Yes: $799	3 PLCR	(800) 582-0524	(708) 808-5000	1367
No	No/No	No/Yes	Yes/Yes/No	84/3	No/No	1 PL	(800) 423-5891	(612) 362-1234	1368
Yes	Yes/Yes	Yes/Yes	Yes/Yes/Yes	83/3	Yes: $105/Yes: $700–$730	3 PL	(800) 225-5627	(513) 445-5000	1369
No	No/No	No/No	Yes/No/Yes	101/4	No/No	1 PL	(800) 538-7506	(510) 661-2220	1370
Yes	Yes/Yes	Yes/Yes	Yes/Yes/Yes	86/2.5	No/Yes: $499	1 PLR	(800) 648-7822	(909) 392-2000	1371
Yes	Yes/Yes	Yes/Yes	Yes/No/No	86/2.5	No/Yes: $595	1 PLR	(800) 788-2878	(818) 369-3698	1372
Yes	Yes/Yes	Yes/Yes	Yes/Yes/No	82/3	No/Yes: $649	3 PLR	(800) 334-3445	(714) 583-3000	1373

N/A = Not applicable   INP = Information not provided   NPA = No price available

New PowerPCs for Notebooks and PDAs

The PowerPC 602 and 603e offer high performance while holding the line on power consumption and cost

TOM THOMPSON

In February, at both the Portable by Design conference and the International Solid-State Circuits Conference, IBM and Motorola announced first silicon for the two newest members of the PowerPC RISC processor line, the PowerPC 602 and the PowerPC 603e. Both processors trace their lineage to the PowerPC 603, which was specifically designed for low-cost and low-power applications (see "RISC Grows up," January 1994 BYTE). Because of this heritage, the PPC 602 and PPC 603e address the same market. However, the IBM and Motorola engineers at the Somerset design center in Austin, Texas, optimized the features of each processor to precisely target specific applications within this market.

The 602, for example, features a die only 50 mm^2 in size as compared to its sire, the 603, whose die measures 85.1 mm^2. Furthermore, the 602 sips 1.2 W at 66 MHz, while the 603 consumes 2.2 W at the same frequency. The engineers achieved this ability by slimming down the

- 1 million transistors
- 144-pin plastic QFP
- 50 mm^2
- uses 1.2 W
- single-precision FPU

capabilities of the 602's FPU, shrinking the size of its on-chip caches and multiplexing the address and data signals on a single bus. These features make the 602 ideal for low-power, small form-factor systems such as high-performance consumer electronics and PDAs (personal digital assistants). The 603e, on the other hand, has larger caches, improved memory management, and operates at 100 MHz. Despite these high-performance features, the designers held the 603e's die size to 98 mm^2, and it consumes only 3 W at 100 MHz. The 603 thus offers capabilities ideal for notebook computers where the emphasis is on long battery life while delivering the utmost in performance. Production of both processors is expected in the second half of 1995.

602: New Chip on the Block

The 602 is 3.3-V part fabricated using a 0.5-micron 4-level-metal static CMOS technology. The die measures 7.07 mm by 7.07 mm and contains 1 million transistors—making it smaller than the 603, which weighs in at

7.4 mm by 11.5 mm and has 1.6 million transistors. (See the table "The Low-Power PowerPC Family" on page 212) Like the 603, the 602 implements a 32-bit version of the 64-bit PowerPC RISC architecture, where the processor supports 32-bit addresses and 64-bit data. However, where the 603 has separate data and address pins, the 602 time-multiplexes the address and data signals on one set of bus pins. This enables the 602 to be housed in a 144-pin plastic QFP (Quad Flat Package), while the 603 uses a 240-pin ceramic QFP. Although this trade-off requires extra bus cycles for data accesses, IBM and Motorola expect the 602's bus to out perform any memory subsystem using 70-nanosecond or slower RAM. That's because most of the 602's memory accesses are made as bursts, which require only that the data's starting address be set up on the first cycle. The 602's small die size and fewer signal lines make it attractive for low-cost applications, where design issues of price and logic-board real estate are paramount.

Other steps were taken to trim logic from the 603 design without compromising the 602's features and performance. The 602 uses a Harvard architecture, with two separate on-chip caches for instructions and data. These caches are 4 KB in size (as opposed to 8 KB per cache on the 603), and each is managed by a separate MMU (memory management unit). The smaller cache size is balanced by the performance of the cache's two-way set-associative organization. Like the 603, these caches support a three-state cache coherency protocol (modified, exclusive, and invalid) that's tailored for a single-processor system design. The 602 has four independent execution units (integer unit, branch processing unit, load/store unit, and floating-point unit). The 602 lacks the system unit found in the 603, because most of its functions were integrated into the integer unit. Since the 602 is expected to operate in single-user environments, such as in PDAs or in embedded graphics-intensive tasks such as a set top box, the designers deemed this consolidation of logic necessary.

Where it was possible and practical, the logic in the remaining execution units has been simplified. For example, the 602's FPU handles only single-precision (32-bit) IEEE-754 standard arithmetic, while the 603 handles both single- and double-precision (64-bit) arithmetic. A logarithmic addition function was added to

- 2.6 million transistors
- 240-pin ceramic QFP or 16-by-16 BGA
- uses 3 W
- pin-compatible with the 603

support speech and handwriting recognition. The designers used a different type of counter to implement the FPU's multiplier array, which reduces its size. The instruction decode logic was simplified so that it fetches and dispatches only one instruction per cycle,

THE LOW-POWER POWERPC FAMILY

PROCESSOR	PPC 602	PPC 603	PPC 603E
Die size	50 mm^2	85.1 mm^2	98 mm^2
Number of transistors	1 million	1.6 million	2.6 million
Number of execution units	4	5	5
Instructions dispatched/cycle	1	2	2
Cache size and type	4 KB, two-way	8 KB, two-way	16 KB, four-way
Power dissipation (typical)	1.2 W @ 66 MHz	2.5 W @ 80 MHz	3 W @ 100 MHz
Performance (integer)	40 SPECint92 @ 66 MHz	75 SPECint92 @ 80 MHz	120 SPECint92 @ 100 MHz
Performance (floating point)	N/A	85 SPECfp92	105 SPECfp92

versus two instructions per cycle on the 603. While this simplification exacts a performance hit, it trimmed the size of the instruction decode and branch prediction logic. Eliminating the sophisticated graphics and string move operations reduced the load/store unit's complexity. However, the remaining move operations were optimized so that they complete within a single cycle. Also, the load/store logic is smarter about accessing data that lies within a double word boundary yet straddles a word boundary: It accesses the cache only once, for a performance boost. Despite these reductions, the 602's MMU offers a fast "protection-only mode" memory-protection mechanism. It uses special bits in the MMU's TLB (Translation Lookaside Buffer) to provide memory-protection functions without the overhead of an address translation. This is ideal for embedded operating systems in which multiple tasks must run in their own memory spaces.

The 602 uses static logic, which preserves the internal state of the caches and execution units when the clock signals to these devices are disabled. The 602 uses the same power-saving modes as those found in the 603: doze, nap, and sleep. The doze mode switches off most of the processor except for the bus snooping logic, which maintains the coherency of the internal caches. The nap mode disables the bus snooping for further power savings. The sleep mode disables the clock to all internal units for maximum power conservation. Even operating at full power, the 602 uses the same DPM (dynamic power management) techniques found in the 603. This is where, depending on processor activity, the DPM logic switches off the clock to idle processor subsystems. For example, if the branch unit has no pending instructions, the DPM logic disables the clock to the branch unit. If a branch instruction is detected in the instruction stream, the DPM reactivates the branch unit. When a tag miss occurs on a cache, the DPM suspends the cache until the bus logic fetches the first beat of data. Finally, the 602 has a PLL (phase-locked loop) logic so that it can be clocked two times or three times the external bus clock rate. This enables the 602 to operate at, say, 66 MHz, while the system runs at 33 MHz, thus reducing overall system power consumption. The 602 is estimated to dissipate 1.2 W during typical 66 MHz operation and 1.4 W maximum. It consumes under 2 mW in sleep mode.

Simulations show that the 602 should perform at about 40 SPECint92 at 66 MHz, compared to 52 SPECint92 for the 603 at the same clock speed. However, the 602 consumes nearly half the power and uses much less board real estate in terms of die size and number of signal lines. Furthermore, most PDA processors don't support floating-point computations, while the 602 offers single-precision arithmetic. This makes it valuable for graphics-intensive applications and for speech and voice-recognition tasks.

603e: Faster and Better

The 603e is, simply put, an enhanced version of the 603. It's a 3.3-V part and uses a 0.5-micron 4-level-metal static CMOS technology. It packs 2.6 million transistors on a die that measures 8.4 mm by 11.67 mm. The processor is available both as a 240-pin ceramic QFP or a 16-by-16 BGA (Ball Grid Array). The first and foremost enhancement is that the 603e supports a faster clock: It can operate at 100 MHz, while the 603 tops out at 80 MHz. The processor also supports a wider range of clock multipliers (1×, 1.5×, 2×, 2.5×, 3×, 3.5×, and 4×), which allows systems designers to hold down power system consumption by using slower clock rates in the system. For example, a system can be clocked at 66 MHz while the 603e using the 1.5 multiplier runs at 100 MHz. Like its predecessor, the 603e has DPM logic that manages the activity of the various subsystems. At 100 MHz, the 603e should typically dissipate 3 W (3.5 W maximum).

The most prominent change to the overall processor design is the large on-chip caches' size: It has two separate 16-KB four-way set associative caches, while the vanilla 603 has two 8-KB two-way set associative caches. The larger caches are possibly in response to Apple's reported difficulties in getting its 680x0 emulator to function on the 603. This emulator uses the 680x0 opcode as an index into a large lookup table that points to the corresponding PowerPC instructions. It was this large table that flooded the 603's caches.

Finally, small improvements were added to the processor design. A key bit was added to a system register that flags a memory protection violation; the 603 software had to examine a combination of bits to determine if such a violation had occurred. This should improve an operating system's response time to a memory fault. Like the 602, most load/store operations now take fewer cycles to execute.

In simulations, the 603e performs 120 SPECint92 and 105 SPECfp92 at 100 MHz. If the shipping part matches these figures, then the 603e easily outguns a Pentium (112 SPECint92 and 82 SPECfp92) running at the same clock speed, all while operating at power levels that make it ideal for notebook computers.

It is interesting that both IBM and Motorola view the low-power market as splitting into two segments: low-cost, single-user PDA or embedded applications, and long-battery-life, high-performance applications. With the 602, 603, and 603e, they have the processors to invade each of those segments. ∎

Tom Thompson is a BYTE senior technical editor at large with a B.S.E.E. degree from Memphis State University. He is an Associate Apple Developer and author of Power Macintosh Programming Starter Kit *(Hayden Books, 1994). Contact him on AppleLink as* T.THOMPSON *or on BIX or the Internet at tom_thompson@bix.com.*

HP-UX 10.0

How HP improved the performance,

reliability, and ease of use of its

flagship PA-RISC operating system

JOHN SONTAG

HP-UX 10.0, the latest operating-system release for Hewlett-Packard's PA-RISC family of servers and workstations, should have just shipped by the time you read this. HP designed it to perform better on SMP (symmetric multiprocessing) systems and under heavy loads.

To increase reliability, the new HP-UX has a journaled file system and the ability to replicate services. It is also easier to install, update, and manage thanks to a new software distribution utility, better system administration tools, and a simpler bundling scheme that rolls features formerly available only on workstations or servers or as add-ons into a single standard package. Finally, it's more standard. Version 10.0 complies with most of the COSE SPEC 1170, the Unix System V release 4.0 file system layout, NFS 4.2 with diskless support, and the Posix real-time interfaces.

Making It Faster

Processors keep getting faster, but overall system performance can't scale accordingly unless I/O throughput does too. Version 10.0 attacks the I/O bottleneck in several ways. It coalesces I/O requests that access sequential areas of a disk into a single I/O operation, which can improve disk throughput by up to 30 percent. Its read-ahead routines

are tuned to predict file access behavior more accurately and exploit I/O coalescence. On SMP systems, disk throughput on a 12-CPU system was improved sixfold by more intelligent management of I/O initiations and interrupts across the set of processors.

Traditionally in Unix, a high-priority process in need of more physical memory can force a lower-priority process to swap to disk. In the case of a large application, the swap can involve hundreds of megabytes, and even with a fast wide SCSI device, the high-priority process might have to wait 5 or 6 seconds for the swap to complete.

In version 10.0, the virtual memory manager achieves much smoother operation under high load by replacing process swap with process deactivation. When physical memory runs low or processes begin to thrash memory, low-priority processes can be deactivated—that is, taken off the run queue. The system then need only displace small clusters of pages when a high-priority process requests memory. Instead of waiting 5 seconds or more, the process waits only about 20 milliseconds—a major improvement in response time.

To speed up systems with a large memory load, version 10.0 adds the `serialize` command (and system call). Consider two concurrent instances of a simulation program, each of which randomly touches a large data array. In earlier versions of HP-UX, these two instances would thrash physical memory because the least recently used and priority page replacement algorithms would not be able to predict which pages to keep in memory for best throughput. The `serialize` command tells the virtual memory system to make a process eligible to be serialized behind other processes. It's analogous to the Unix `nice` command. When there is plenty of memory, `serialize` has no effect. But when memory runs low, the system runs serialized processes one at a time. Once the highest-priority serialized processes run to completion, the lower-priority processes can complete.

In some test cases where two instances touch large arrays of data randomly, this technique cuts clock time by a factor of eight. As with the `nice` command, you can return a serialized process to normal priority.

With Process Resource Manager, an alternate process scheduler, the system administrator can create groups of users and guaran-

High Availability with MC/ServiceGuard

A — Node 1

Processes:
app_process_1
app_process_2
middleware_1
middleware_2
Disks:
lvol_data1
lvol_data2
Network:
IP 56.23.101.44

Package A
Package B
Package C

B — Node 2
Package D
Package E
Package A

Node 1
Package A
Package B
Package C

Node 3
Package F
Package B
Package G

Node 4
Package C
Package H
Package I

(A) Administrators define packages of resources that applications need to run, and they assign packages to servers.
(B) When a server fails, MC/ServiceGuard migrates packages to backup servers in the cluster.

tee each group a minimum share of the total CPU time. Mission-critical applications must respond to users in a hurry. This technique ensures that background activities don't get in the way, even if they run at a higher priority than the mission-critical software.

Making It More Reliable

LVM (Logical Volume Manager), which came from OSF/1, manages collections of disk drives. It partitions drives, mirrors data for redundancy, and stripes data across multiple disks for higher performance. HP-UX 10.0 upgrades these features to improve the resiliency of mass-storage subsystems. When you're running a 24x7 (24 hours a day, 7 days a week) operation, there's never a convenient time to do a backup. Now you can remove a drive from a mirrored pair to enable off-line backup from any node with no interruption of service. LVM also lets you take one copy of a data set spread across many disks off-line and back it up. The backup utility can access a frozen set of data and operate at high speed. When backup is complete, the disks are brought back on-line and synchronized with the live system data. LVM exploits RAID disks with multiple controllers, offering an automatic fail-over capability.

Version 10.0 further protects data integrity with a journaled file system, VxFS (the Veritas file system) (see "The Great Little File System," February BYTE). Compared to the BSD 4.2 HFS and NFS, VxFS has superior data integrity, recovery, and performance. It maintains an intent log of uncommitted meta-data transactions. If the system crashes, recovery is a simple process of reading the intent log and applying or backing out changes. With add-on products, it is possible to resize and reorganize file systems on-line, control caching options, and use the intent log for fast, synchronous writes. All these features add up to a file system that is much more resilient across system failures and, potentially, much faster.

Resilience to memory faults becomes increasingly important as system memories grow larger and denser. In version 10.0, the diagnostic system and the operating system can mark bad pages and then avoid using them, thereby preventing system panics. If a page shows two occurrences of recoverable errors at the same address or one unrecoverable error, it's removed from service. Information about bad pages resides in a nonvolatile RAM, where it survives across system boots.

Ultimately, of course, reliability means keeping applications up and running no matter what. To that end, version 10.0 offers MC/ServiceGuard, a facility for mutual backup of services across clusters of up to four servers. When a protected service (or the system supporting it) fails, MC/ServiceGuard resurrects it on another system in the cluster. Applications are made highly available, without having to be rewritten, by means of packaging.

A package defines the set of resources an application needs to run, including disks and network resources. When a system, network connection, or application fails, a clusterwide monitor notices the service interruption and launches a package on a backup system. Using multiple disk connections, the backup system can commandeer the failed system's disk drives, and its networking interfaces can adopt the failed system's IP address.

When a mirrored disk or a redundant network interface fails, repair can occur in under 10 seconds. If an entire system fails, users will be able to continue in 1 to 2 minutes, once application recovery is complete. Depending on the application, you might lose some data entry, or you might even have to reenter the application—this isn't nonstop computing. But it's an extremely cost-effective way to have servers back up each other. It offers peak performance when all is well. When a failure occurs, service remains available with some degradation of performance because one system now must do the work of two.

Making It Easier to Use

HP-UX's system management tools share a common interface thanks to the OBAM (Object-Action Manager), which encapsulates diverse disciplines, including X Window System, Motif, international and CDE (Common Desktop Environment) support, on-line help, regression testing, and character-based terminal support. Using the GUI-based SAM (System Administration Manager), an HP-UX administrator configures and manages auditing and security, backup and recovery, disks and file systems, diskless cluster configuration, the kernel and devices, networks, peripherals, printers, processes, and user and group accounts. New with version 10.0 is a major reorganization of SAM, with emphasis on typical administrative tasks. The administrator also delegates such tasks to other users—with appropriate security restrictions—and adds user-defined utilities to the SAM menus.

Version 10.0's Software Distributor, or SD-UX, includes tools that package, distribute, and manage applications and operating-system software, as well as data. With SD-UX, users can pull software off the network and install it locally. Using an add-on product, administrators' networks will be able to push software to any of the nodes in the system. SD-UX runs on top of a DCE (Distributed Computing Environment) RPC (remote procedure call) and will exploit a secure RPC, as well as DCE authorization, authentication, and directory services where available. SD-UX also lets customers define bundles made of products and partial products, install them on their systems, interrogate systems to determine what is installed, and remove software from systems.

Making It More Standard

HP-UX 10.0 adds support for many industry standards. It complies with most of the COSE SPEC 1170 and lacks only System V signals and internationalized curses, both of which are due later this year. The HP-UX real-time scheduler, available since 1986, complies with Posix 1003.1B, which defines interfaces to a real-time scheduler and a set of high-resolution timers.

HP-UX 10.0 also converts to the standard Unix SVR4 file system layout, so SVR4-oriented users can find files in the directories where they're traditionally kept. (Links to the HP-UX 9.0–style directories ensure compatibility with the prior HP-UX tradition.)

HP-UX 10.0 bundles the DCE client. The base technology, from OSF 1.0.3, has been upgraded to include the security and RPC features of OSF 1.1. The disk footprint for the DCE client shrinks by 75 percent, and the RPC code has been tuned for about a 30 percent performance boost. Version 10.0 also bundles the Streams architecture with the base HP-UX product.

With NFS 4.2 support, version 10.0 enables diskless systems to be served using the NFS diskless protocol. Initially for HP systems only, this capability will later extend to other vendors' systems, too. Version 10.0 also offers full support for the 4-byte EUC (Extended Unix Code) code sets, so programmers can internationalize their applications.

HP considers HP-UX a unified product—one that is equally at home on a small uniprocessor desktop machine or a 12-way SMP superserver in the data center. ∎

John Sontag is the chief architect for HP-UX at Hewlett-Packard and has been working with HP-UX on PA-RISC systems since 1983. You can reach him on the Internet at sontag@cup.hp.com or on BIX c/o "editors."

The Software Stopwatch

It doesn't take special hardware

to achieve timing resolution

in the microseconds on the PC

or the Mac. You just have to know

where to look.

RICK GREHAN

O ptimization is often the name of the programming game. But before you can optimize, you have to find that which is not optimal. Bill Atkinson, a member of the original Mac design team once said: "Optimization without measuring is wasted time. Find out where the application's really spending time and go whump on that code" (February 1984 BYTE, page 76).

So, to optimize, you must whump, and before you whump, you must measure. But measurement may not be an easy task with today's processors running 200 MHz and higher. A person with a Rolodex won't even begin to cut it. Some type of hardware/software assistance would be nice, but how do you safely get at all those timers and hardware things that the operating system works so hard to insulate you from?

Time on the PC

At the heart of the old IBM PC—and still beating in PC clones today—is a hardware timer that issues an interrupt about once every 55 milliseconds. This works out to about 18.21 clock ticks per second and is the basis for the time-of-day clock on most PCs.

It's also the reason "generic" time functions on a PC may be inadequate for high-accuracy timing. Suppose you've written a program whose performance is highly dependent on a sorting routine. You have one of two choices as to what sort algorithm to put at the core of that routine. In your tests, the first sort algorithm takes a twentieth of a second; the second takes only a fortieth of a second (not unheard of in these days of 90- and 100-MHz Pentiums).

If you simply test the candidate algorithms by running an iteration of each and timing the duration, a DOS clock-based timer will report that they both take the same amount of time. The problem is one of resolution; a clock that ticks only 18 times a second can't "see" any events shorter than that.

Brief Tangent

Before we discuss fixing this problem, stop for a moment for a warning. The clock() function in many DOS C compilers will appear to return results with resolutions to the hundredths or thousandths of a second. The warning: This is only partly true. The true part is that the number returned by clock() has the proper dimension. The not-so-true part is the fact that the clock() routine doesn't advance one unit per tick. It lurches forward in time, skipping multiple milliseconds at a tick.

For example, you can write a short C program that repeatedly calls the clock() routine, displaying a value only when there's a change. (Try this with your favorite DOS compiler.) We tried this with one DOS compiler and the value of clock() advanced by 5 or 6 ticks. That compiler's time.h header file told us that the clock() function presumes 100 ticks per second. Sure enough, 55 ms works out to between five- and six-hundredths of a second, given that clock() will exhibit a 5-ms "jitter."

Higher Resolution

There is a route to higher resolution on a PC running DOS. It appeared in the January 1987 BYTE in Byron Sheppard's article "High-Performance Software Analysis on the IBM PC." The article included the source code

---

A short assembly language routine to fetch the entry point of the VTD. It returns the address in DX:AX.

```
getapi_
  proc near
  mov   ax,1684h   ;Subcode to return API
  mov   bx,05h     ;VTD ID
  xor   di,di      ;Clear ES:DI
  mov   es,di
  int   2fh        ;Call interrupt
  mov   dx,es      ;Return results
  mov   ax,di
  ret
getapi_ endp
```

Call the VTD. The entry point is passed to the routine in the DX:AX register pair. Note that the routine fakes a CALL using a far return. Also, this routine presumes tarray is a two-element array of double words.

```
gettick_
  proc near
  push cs
  mov   bx,offset retspot
  push bx
  push dx          ;Push VTD segment
  push ax          ;Push VTD offset
  mov   ax,100h    ;Function code
  retf             ;Make the far call
retspot:
  mov   dword ptr _tarray,eax
  mov   dword ptr _tarray+4,edx
  ret
gettick_ endp
```

listing for assembly language routines that could easily be modified for calling from high-level language programs.

Sheppard's approach involved reprogramming the PC's timer 0, which turns out to be the timer that generates the 55-ms clock ticks. It also turns out that timer 0 is really ticking away with 840-nanosecond pulses; the BIOS programs it to count 65,536 pulses before generating the interrupt. In a nutshell, Sheppard's routine reads the count in the timer, which—with some math—you can use to determine the number of 840-ns ticks since the last timer interrupt. This gives you better-than-microsecond accuracy. If you want to check out Sheppard's listing, you can dig up your back issues of BYTE or download the file JAN87.ARC from the listings/frombyte87 area on BIX.

Virtual Timers

Unfortunately, using a hardware timer isn't always an available option. Within operating systems (e.g., Windows, which virtualizes hardware), directly accessing hardware can yield bizarre results. We tried Sheppard's high-resolution timer code from within a Windows program: Sometimes it worked, and sometimes it didn't.

Fortunately, Windows has a kind of back door into a timer that is as good as Sheppard's. Specifically, Window's VTD (virtual timer device) provides access to a timer with a resolution of—guess what—840 ns. You can get to it using two assembly language routines lashed to your main code.

One routine calls INT 2Fh, which is a kind of clearinghouse interrupt that returns the API entry point for all the device drivers Windows knows of. You just plug a virtual device ID into the BX register (the ID of the VTD is 05) and call the interrupt; ES:DI returns holding the entry-point address in segment:offset form (or all zeros if the VxD has no entry point).

The other routine actually calls the VTD. Plug a 100h in the AX register (this is the function code) and then call the entry point that the first routine returned. The VTD will return the number of 840-ns ticks because Windows was started in EDX:EAX. The listing on page 215 shows both routines.

Time on the Mac

Suppose you're a Mac developer, and you want to do some high-resolution timing. There are no hardware timers to reprogram here. You first might try to use the timer global variable residing at 016Ah (referred to as LMGetTicks in the standard header files for the Mac). This location is updated 60 times a second and holds the number of ticks since the Mac was started. It is at least better than the DOS clock.

But the revised time manager—provided with the Mac OS 6.0.3 and higher—can provide timing with accuracy to 20 microseconds. (The Mac OS 7.0 has the extended time manager, which does everything the revised time manager can and more.)

The time manager's real job is to schedule tasks to run at predetermined future times. This lets you set up tasks that run at regular intervals, an important feature for multimedia and real-time activities that need routines run at precise intervals.

What we want, however, is the ability to measure durations of time. With the above description of the time manager, it seems that the logical approach to this is to create a routine that wakes up every millisecond or so and updates a global variable; sort of a high-resolution form of LMGetTicks.

But Apple has built into the time manager a way to measure *without* submitting a task for the manager to run. Sounds weird, but it's true. The time manager keeps a list of all the tasks scheduled to run on a queue, a linked list of data structures. Each queue member contains a pointer to the task that you want awakened in the future, as well as the delay (i.e., how far into the future the manager must go before waking up the task). You can specify the delay in milliseconds (for a long delay and not much accuracy) or microseconds (for the reverse). The time manager knows whether you want milliseconds or microseconds by examining the sign of the delay field: A negative value indicates microseconds, and positive indicates milliseconds. You put an item onto the queue using the InsTime routine (after building the appropriate data structure, of course) and "arm" it with a call to PrimeTime. This latter function is the one that tells the time manager when to run the task. You remove an item from the queue with a call to RmvTime.

If you place an item on the queue that has a NIL value in the task pointer field, the time manager never starts the task (which makes sense, because you've basically told the time manager that there is no task to start). But if you issue a PrimeTime call on such a queue element, the time manager keeps track of how much time is *left* were the hypothetical task started (it places this duration in a field of the queue element called tmCount). You then call RmvTime and examine tmCount to determine the time left, from which you can compute the time that has elapsed since the call to PrimeTime. The rest is obvious: Bracket the code you want to time between a call to PrimeTime and RmvTime, and you have a software stopwatch (see the listing below).

Watch the Watch

Getting at high-resolution timers doesn't have to be a programming nightmare. Of course, you've got to do some exploring. We had to dig through *Inside Macintosh* and various Microsoft developer CDs—not a carefree jaunt. But it paid off. Now we've got stopwatches accurate into the microseconds.

Whump away. ■

Rick Grehan is the technical director of the BYTE Lab. He has a B.S. in physics and applied mathematics and an M.S. in mathematics/computer science. He can be reached on the Internet or BIX at rick–g@bix.com.

Using the time manager on the Mac as a software stopwatch. This code can time events up to 10 minutes. Note the calculation of ohead *, which factors out the overhead of the time manager routine calls.*

```
struct TMTask myTMT;
long delay, ohead, rslt;

/* Clear TMTask struct */
memset((void *)&myTMT,0,sizeof(TMTask);
delay=100*1000*1000;    /* 10 minute delay */

/* Put task on queue */
InsTime((QElemPtr)&myTMT);

/* Calculate overhead */
PrimeTime((QElemPtr)&myTMT,-delay);
RmvTime((QElemPtr)&myTMT);
ohead=delay+myTMTask.tmCount;

/* Time something */
InsTime((QElemPtr)&myTMT);
PrimeTime((QElemPtr)&myTMT,-delay);
..insert stuff to be timed here...
RmvTime((QElemPtr)&myTMT);

/* rslt has duration in microsecs */
rslt=delay+myTMT.tmCount-ohead;
```

Create More IP Addresses

The transition to a new IP address

format doesn't have to be painful

TIM WINSTON

The explosive growth in demand for IP addresses has led to the development of a new IP address format called IPng (IP next generation). Eventually, all internetworking equipment will run the new format because of its expanded addressing capability and also because of other new features incorporated into IPng that simplify packet routing and handling.

The challenge for a corporate network manager is how and when to make the transition from their current IP (IP version 4) network to one based on IPng. Fortunately, there are several ways to make the migration.

Companies have the option of upgrading routers to start, or they can mix IP host and router upgrades. All the while, they have to maintain compatibility with existing IP internetworks that may have different migration strategies.

Why Change?

TCP/IP's growth has prompted the need for a new version of IP addressing. The Internet, the largest IP internetwork, is comprised of over 40,000 networks. This number has been doubling about every 12 months.

The most apparent threat is depletion of the 2^{32} available IP 4 addresses that represent over 4 billion nodes. Although it seems like that's enough addresses to go around, there's a potential shortage because of the way addresses are assigned.

IP addresses are allocated in chunks per network, not by the individual device. The minimum allocation is 256 addresses. Yet, a network backbone may use only two of these (one for a router and one for a host computer). This hierarchical assignment is necessary for routing, but it reduces the practical number of available addresses.

A more pressing problem is that routers, especially those on the Internet, are rapidly being overwhelmed. Backbone routers on the Internet must effectively keep a table of every network address. A router must use this table to look up the destination address of each packet it receives before it can forward the packet. Before it receives the next packet, the router must complete the entire process (address lookup and forwarding).

While new protocols for sharing route information and new hardware are in continual development, the growth of internetworks may outpace such developments. Thus, the potential shortage of IP addresses and the burden on routers require that a new version of IP be deployed within three to seven years.

The IETF (Internet Engineering Task Force) and the Internet community began looking at proposals for the IPng in early 1992. As the debate progressed, proposals merged and evolved until the IPng area directors made a recommendation in September 1994. The result is that IPng (technically, IP version 6) has several key features that distinguish it from IP 4:

• IP 6 offers greatly expanded routing and addressing capabilities. IP 6 addresses are 128 bits, compared to 32-bit IP 4 addresses. While this obviously provides astronomically more addresses, it mainly supports more levels of addressing hierarchy to simplify routing.
• IP 6 has simpler packet headers. IP 4 headers include options that are not necessary for routing, but routers must process them. IP 6 has a fixed-size header containing only information needed to route the packet. Options that only the final recipient needs are put in subheaders and don't have to be processed for the router to send the packet to the appropriate output port.
• IP 6 lets you label packets for special handling, such as "real-time" service.

IPng Handles Mixed Addressing

IP 4/IP 6 host · IP 4 host · IP 4 host · IP 6 header / IP 4 header / Data · IP 4 host

IP 4 header / Data

IP 4/IP 6 router · IP 4 router · IP 4/IP 6 router · IP 4/IP 6 router

IP 4 header / Data · IP 4 header / Data · IP 4 header / Data · IP 4 header / Data

The flexibility designed into IPng ensures compatibility even when there's a mix of current and new IP addressing formats. Here are just two examples of the types of mixed environments that can be supported.

Other protocols, most notably OSI (Open Systems Interconnection), have solved some of these problems that face TCP/IP. But none had an adequate transition plan. For example, an OSI router can't support a computer running only TCP/IP. Before the transition to an OSI network is complete, all devices must be upgraded.

The IETF avoided this problem by ensuring a smooth transition from IP 4 to IP 6. The IETF's primary goal for its transition plan, called SIT (Simple Internet Transition), is that it be easy. If it is difficult, people won't do it, and IP 6 will fail.

One key element in meeting this goal is that incremental deployment is possible in a transition to IP 6 networks. Any router, computer, or other device can be upgraded without requiring anything else to be upgraded simultaneously (see the figure "IPng Handles Mixed Addressing" on page 217). This means that IP 6 support can be added in the regular maintenance cycle of each device.

Furthermore, devices do not need to be upgraded in any specific order. Routers can be upgraded at any time, regardless of the state of other routers or hosts (computers and other end nodes). The exception is that hosts need an upgraded DNS (Domain Name System) to resolve names for IP 6 devices. Existing IP addresses that are compatible with the new addressing structure also make the transition easier. IP 4 addresses map to IP 6 addresses.

If you've been allocated IP addresses, you can use them according to your plans while upgrading your routers and hosts to IP 6 software. By the time you run out of your allocation and can't get another IP 4 allocation, the software transition should be complete. Until then, you won't need an address deployment plan.

While IP 6 supports IP 4, it does not extend it. As IP 4 depletion nears, hosts will begin using addresses outside of the IP 4 "network." Any device that does not have its software upgraded (e.g., a device that is IP 4 only) will not be able to interoperate with these "IP 6–only" devices.

When devices are upgraded, the IP 4 "network" is divided into two parts, IP 4–only (all TCP/IP hosts), and IP 6/IP 4 (hosts that do both). The IP 4–only addresses are represented by their 32-bit address padded out to 128 bits with zeros. For example, 0:0:0:0:0:0:192.128.3.24 (numbers followed by colons represent 16 bits, 8 bits if followed by periods. And IP 6 addresses are represented by eight 16-bit numbers separated by colons). When an IP 6/IP 4 node sees this type of address, it knows it must use IP 4 to communicate.

Addresses for nodes that handle both IP versions look the same with one exception—the 16 bits preceding the IP 4 address. For example, the address would look something like 0:0:0:0:0:FFFF:192.128.3.24. This identifies the node to other IP 6 devices as an IP 6–capable node, but it also has a unique IP 4 address to interoperate with IP 4–only devices.

As the transition progresses, enough devices will be upgraded to IP 6 that it will become practical to assign addresses outside of the IP 4 "network." The addresses that use the higher order bits will not be inherently interoperable with IP 4-only nodes. However, any IP 6/IP 4 node will interoperate, even if it is assigned an IP 4–compatible address.

There has been discussion of a mechanism to map IP 6 addresses onto the IP 4 network. This would let a local network stay IP 4 only, yet interoperate with remote IP 6–only devices.

Implementation Strategies

Taking this level of compatibility designed into IP 6 and applying it to internetworking equipment is straightforward. For most companies, routers will almost certainly be the first devices to gain IP 6 capability. Routers use addresses, not names, so they are not dependent on any other IP 6 support. All existing router protocols work with IP 6, so common router-to-router links will probably be the first IP 6 networks (tunneling IP 4 across IP 6 networks).

Because computers and other end nodes need to resolve the mapping of names to IP 6 addresses to establish an IP 6 connection, DNSes need to be upgraded to support an additional record type. Thus, new records need to be added for each device as it becomes IP 6–capable.

One of the new features of IP 6 is support for auto-configuration (plug and play) and mobile devices. This will necessitate changes to make DNS dynamic so that roaming devices can update their location on the network based on their name. Authentication is a key part of mobile support. It allows a specific device to be positively identified before updating its network location.

Network management software will also need to support the (as-yet-undefined) MIB (Management Information Base) for IP 6. Support for monitoring authentication, auto-configured devices, and mobile devices will also have to be added.

Servers should have simultaneous support for IP 6 and IP 4. They will be used by devices with different transition schedules and priorities. Thus, they should be as accessible as possible, both locally and remotely (remember, a remote device may not be upgraded on the same schedule). Upgrading servers also can move heavy server-to-server traffic to IP 6, helping the routers and improving performance.

Desktops and other end nodes vary greatly in their use and priority of upgrade (and the availability of upgrades). General-purpose computers, such as PCs and workstations, will probably have dual-protocol support available about the same time as servers. IP 6/IP 4 capability is important to give them access to remote hosts and stay independent of each host's upgrade schedule.

As printers and terminals wear out, their replacements will probably support IP 6, but as long as the computers that use them are dual protocol, these devices may never have to be upgraded.

Rules to Live By

When faced with a transition from IP 4 to IP 6, the main thing to remember is don't panic. IP 4 will exist for a long time. The move to IP 6 will prolong the viability of IP 4 networks by freeing up addresses and increasing the efficiency of routers.

For the most part, network software upgrades will "automatically" add IP 6 capability as vendors add IP 6 support to their products to stay competitive. This will be part of the normal upgrade and maintenance procedures. And existing address assignments will work automatically.

Upgrades should be available before IP 4 addresses are depleted. By the time you need to deploy new addresses, auto-configuration, which is planned as part of IP 6, should be defined and available. Auto-configuration assigns addresses to devices as needed. It should also permit centralized address management.

As devices are upgraded to IP 6, simply add a new record to the DNS so that other upgraded devices will be able to use an IP 6 connection to reach them. Taking these steps and the backward compatibility of IP 6 into account, should allow a smooth transition from IP 4 to IP 6. ■

Tim Winston has worked extensively in the design of large corporate networks. He is a development group manager at WRQ, a PC connectivity company with headquarters in Seattle, Washington. He can be reached on the Internet at timw@wrq.com or on BIX c/o "editors."

JERRY POURNELLE

Orchids and Onions: Part 1

It's time for the annual Chaos Manor User's Choice Awards and the annual Orchid and Onion Parade. I do these now because, unlike the rest of the publishing world that believes a year ends in early November when the January issue goes to bed, I insist that a year isn't over until the end of December. Many products don't come out until the November Comdex, and while I could frantically try to review those, I'd never be able to use them before my deadline for the January issue.

Ground rules: these are my awards. I consult with the other BYTE editors, but I decide. A User's Choice Award means: I've done enough with the product to become familiar with it. I've probably used the product. I certainly like it, and I recommend it as good enough. Nowadays, I seldom pretend to know what is the "best" in any category. There's just too much out there.

When I say I use something, I generally mean it. Despite the many piles of

Jerry presents his list of last year's best and worst computer products

software, Chaos Manor is a small business. I do research, build and test mathematical models of dynamic processes, write programs, write and market science fact and fiction, buy and sell stories for my anthologies, hire consultants, keep books, pay my bills, and do my taxes. I'm always looking for ways to use computers to make my life easier, and that's generally what determines a User's Choice Award.

The User's Choice Award for Operating Systems
goes to OS/2 Warp, but a warning goes with it.

OS/2 Warp works just fine for people like my wife. Roberta uses Desqview and DOS programs at present. She really needs Windows, but she doesn't much relish learning it. The only multitasking she does is running communications in the background while she uses a word processor. OS/2 Warp would let her run a few Windows applications without knowing much about Windows and would be easier to learn than Windows. It's also greatly superior to both Desqview and Windows for task switching among DOS programs.

It would give her good access to the Internet. OS/2 Warp comes with pretty good access tools.

The mail program has a miserable user interface and sometimes doesn't properly access the Internet; but the WWW (World Wide Web) browser is pretty cool, anonymous file transfer works well, the news browser isn't bad, and they all work smoothly in the background.

Moreover, OS/2 Warp is stable. Once it's set up properly, it's extremely easy to use. You can nest folders and program groups, set which ones open on start-up and which don't, and, in general, set up buttons to do routine stuff for you. Once all this is done, OS/2 Warp will work reliably and efficiently.

The catch is setting it up.

That doesn't have to be difficult. Some people will have no trouble at all, particularly those who have a CD-ROM drive that OS/2 recognizes; but even the floppy drive version can be installed easily, albeit tediously, if you have the right equipment, even if you don't know much about what you're doing. The lucky ones may then find it simple to access the Internet and join those crowds of smiling people seen in the OS/2 Warp ads as they go network surfing and have a great time of it.

Others won't be so lucky. Some will have

Pournelle

equipment incompatibilities. Others—well, I don't know, they just may not hold their mouths right. For whatever reason, installation will be sheer hell. Favorite software packages won't work. The system will crash mysteriously. Perhaps Windows won't work, even though it's your own Windows. Unlike OS/2 2.11, Warp doesn't come with Windows. You'll need to have your own, and, if you're smart, you'll have it working properly under DOS before attempting Warp. Windows generally works fine within Warp—but then everything else generally works just fine, too.

Alas, in a sizable minority of cases, enough goes wrong that unsophisticated users will simply give up on OS/2 Warp. For the rest, those for whom things go easily or who can find knowledgeable people to help them set up OS/2 Warp, the result can be very much worthwhile. Warp is fast. If you normally use your machine for the same thing day after day, you'll probably be better off with OS/2 Warp once you get it up and running.

Another group who will appreciate OS/2 Warp is experimentalists, people who like to play around with their computers to see what will happen. If you're looking for tricks and shortcuts, and you have time for messing about with your system, OS/2 Warp may be ideal. It's far more customizable than Windows, and it prints a lot faster, too, even from Windows programs. When I have to print a 600-page novel, I always use the network to transfer the file to the OS/2 system before letting Word 6.0c for Windows have at it.

You'll notice that I said I transfer the file to the OS/2 machine, meaning I don't use that machine for writing, which is true. My primary machine remains Big Cheetah, which is a 486DX2/33 running DOS 5.3, BOOTCON 2.11, QEMM 7.5, Windows for Workgroups, and a DAT (digital audiotape) backup system running under Palindrome's Network Archivist backup software. This is partly due to inertia; but it's also an extraordinarily stable and reliable system.

More to the point, it's very *forgiving*. That's important, because I am forever installing and removing software on this machine. I have to—the only way I can *use* something is to have it on the machine I do most of my work on. Alas, some of the software I try doesn't work very well, and it's important that I be able to remove it easily and, in case of a crash, recover quickly. Big Cheetah with the Palindrome

DAT fills that bill nicely.

OS/2 Warp doesn't. Once Warp is tuned up and running properly, it's stable and reliable, but my experiences using it as an experimental system haven't been very pleasant. For one thing, you sure don't want to reboot often if you're using OS/2 Warp (or OS/2 2.11 for that matter). It takes longer to reboot Warp on a Pentium system than it does to bring up DOS and Windows on my 486—and if you install as much bad software as I do, you'll find yourself rebooting often.

Of course, most of you don't have to install a lot of bad software, if only because you have me to do that for you; and I do use the Warp machine (an IBM ValuePoint Pentium system) as my primary communications system. It does my MCI mail, BIX, GEnie, and my increasingly complicated Internet transactions. Most DOS games play just fine under Warp; and Warp has a dual-boot feature that makes it very easy to boot up DOS for those games that go directly to the hardware and want the entire system.

OS/2 Warp isn't a perfect product, but for many of us, it's more than good enough; and with those warnings, it gets my User's Choice Award for Operating Systems.

I've discussed BOOTCON in several previous columns. It's a program that lets you reconfigure your DOS system each time you boot up. You can have special configurations for games. Some games insist you have a memory manager like QEMM or EMM386.SYS; others, like Wing Commander 3, are happier if you have no memory manager at all, other than their own. Some games want one sound-card configuration, and some want another. With BOOTCON, you can experiment with new programs without fear of mucking things up beyond repair.

BOOTCON works just fine with OS/2's dual-boot feature. That is, when you boot up DOS, you'll see BOOTCON; when you boot up OS/2, the system will never know it's there. They've recently improved BOOTCON yet again, and it's one of the first things I install when I get a new DOS system. I can't live without it, and if you play games or experiment with DOS configurations, you can't either. BOOTCON gets a User's Choice Award for Utilities.

The most powerful machine at Chaos Manor is Pentafluge, a Pentium system that's built from a PC Power & Cooling tower case, a

Smart Connectivity does more than simply get you from here to there.

It connects you simply, easily, and securely. It offers you endless solutions. Smart Connectivity gives you the power to quickly go where you need to go. It gives you the freedom to navigate the information age. Smart Connectivity works seamlessly and keeps pace with your changing needs. Smart Connectivity offers solutions for success.

SMARTERM® is Smart Connectivity.

SmarTerm offers smart emulation and TCP/IP connectivity choices for Windows, 32-bit Windows, and DOS. SmarTerm is the most precise terminal emulation for UNIX, VAX/VMS, MV, and AViiON hosts and includes TCP/IP (a Windows Sockets DLL) and LAT protocol stacks. It works in every network environment—both Ethernet or Token Ring. SmarTerm is easy to use and support with its powerful automation tools and top-notch Technical Support Team.

Evaluate SmarTerm's Smart Connectivity. Call Persoft today at 1-800-368-5283.

"PC Week Labs recommends SmarTerm..."
-*Michael Blakely, PC Week*

Circle 90 on Inquiry Card.

persoft®
CONNECTIVITY SOLUTIONS
DOS • Windows • Ethernet • Token Ring

Turbo-Cool 300 tower-model power supply, and a Pentacool fan; a Micronics Computers M5Pi motherboard; an ATI Technologies Graphics Pro Turbo Mach 64 graphics card; a Distributed Processing Technology SmartCache III SCSI Host Adapter; and a Digital Equipment DEC 3107 1.05-GB digital hard drive (see my September 1994 column).

We had some software problems getting it set up, and you may recall the one hardware problem when I got peanut shells in the keyboard, resulting in my thinking there was a motherboard failure when there wasn't. Otherwise, it has run quietly and efficiently. Even in midsummer the Pentium chip was plenty cool, mostly because PC Power & Cooling gave a lot of thought to the airflow in the box.

Over the years, I have had a number of machines with power supplies from PC Power & Cooling, and I have not been unhappy at all. It's about time the company got some well-deserved recognition: PC Power & Cooling wins the User's Choice Award for Power Supplies.

Microsoft has made significant improvements to DOS, particularly to the EMM386.SYS memory manager, which has an excellent memory test and does a good job of getting around the 1-MB DOS memory problem. However, as quickly as Microsoft improved it, Quarterdeck moved faster. QEMM 7.5 is more efficient and easier to use than the memory manager you get with DOS. You also get Manifest, a program that tells you quite a lot about how your memory is being used, and a densely written manual that may not be easy to read but is well worth studying. The manual is as good a text on memory and memory management as you'll find anywhere.

I've used Quarterdeck products about as long as the company has existed, and while I sometimes complain about one feature or another, I'd be lost without QEMM. QEMM 7.5 gets the User's Choice Award for Memory Management Utilities.

Quarterdeck also has a little Windows shell program that I like a lot. It's a Program Manager replacement called Side-Bar. While it takes a little getting used to, it's worth the effort to learn it. It's also very OS/2-like. SideBar was my favorite Windows utility of the year, and gets a User's Choice Award in that category.

There are two ways to set up OS/2 Warp. One is with HPFS (High Performance File System); the other is with plain old FAT (file allocation table) of the kind used by DOS. I think some small part of your system has to be FAT, so you can use DOS

and Windows. The rest can be HPFS if you like. In my case, I chose to stay with all FAT, because it makes it easy to get back to DOS and Windows. Warp has a dual-boot command that will do just that.

There are advantages to HPFS, particularly when it comes to installing and managing a network. On the other hand, you cannot easily transfer files from the HPFS partition to the FAT partition. Also, if you have a mostly DOS and Windows network (as I do), communications become extremely difficult; this was another reason I stayed with FAT all the way.

This morning I exited OS/2 Warp to boot up DOS. While I had Valiant (the ValuePoint Pentium P5/60) in DOS, I used Golden Bow Systems' Vopt to defrag the system's disk. As you probably know, disk operating systems tend to write files as they find space for them. After a while, the files are splintered into lots of pieces scattered all over the disk. This is known as fragmentation, and it can greatly increase the time needed for file access.

I hadn't actually expected there to be any need for defragging. Valiant has a 500-MB hard drive with fewer than 200 MB of files. Fragmentation generally takes place on full disks. Not this time, though. Vopt shows you a map of the disk before it begins. I was astonished: about half the disk was empty. The other half was jampacked with thoroughly fragmented files. Anyway, in about 4 minutes, Vopt had the disk neat and orderly again.

OS/2's HPFS doesn't need any disk defragmentation, and there are mixed opinions about whether you want to use a defragger on OS/2 FAT disks. I did, and it doesn't seem to have done any harm; but you're probably better off using Vopt on DOS systems, with or without Windows. Of course, Vopt is a DOS program. It won't work in a DOS window, because Windows keeps some files open, and Vopt needs to get at them all. I use Vopt every week or so when I shut a Windows machine down and restart it. It takes only a few minutes, and sometimes it speeds up disk operations something wonderful.

Golden Bow's Vopt isn't a glamorous product. It just works. I've used it for about five years now. It has always been the safest disk defragger I know of, and it deserves a User's Choice Award.

If you program in BASIC, you need Crescent Software's programming tools. I make no

secret of my belief—call it a prejudice if you like—that compiled Microsoft Professional Basic and Visual Basic are the fastest and most reliable ways to get complex DOS and Windows applications up and running. Once you have them running, you may or may not want to look into ways to optimize performance.

Crescent Software's programming tools come in two main varieties. Tested, reliable routines written in BASIC let you add menus, mouse support, and other functions. Routines written in assembly language give you a capability not found in BASIC or greatly speed up functions that are already in BASIC.

If you do program development and don't use BASIC to get things running, you may want to rethink your strategy; if you do use BASIC, you definitely need Crescent Software tools. There are a whole bunch of them, and rather than choosing one, I'm giving a User's Choice Award in Programming Utilities to Crescent Software's whole line.

Zenith pretty well invented practical laptop computers and for a while thereafter dominated the field. The competition caught up and ran past Zenith a few years ago, probably as a result of overmanagement from Groupe Bull. Now the owners have wisely given engineering and development control back to the geniuses at St. Jo, and the result has been spectacular.

When the USS *Tripoli* was deployed to the Persian Gulf last spring, my son Phillip took a Zenith Z-Noteflex color laptop. Like all new Zenith laptops, it came with Windows and power management utilities.

I had promised Zenith a major torture test of the machine. It got that. For a while, the temperature aboard the *Tripoli* was high enough to melt paint off interior bulkheads. They went through typhoons. That machine got wet, dropped, overheated, and generally abused.

Two problems developed. Phillip's model has a trackball that attaches externally to the front of the machine. This keeps the laptop smaller, but the detachable trackball is subject to damage; the first one stopped working before they got to sea. The Zenith service department replaced it under warranty, no questions asked. Phillip reports that the replacement worked throughout the deployment but needed frequent cleaning. Zenith portables come in a slightly larger model with the trackball built in, which makes for a considerably

Pournelle

more rugged machine.

Ship's power on a helicopter carrier is unreliable at best. Phillip had a surge suppressor, but the constant power spikes fried it, after which a surge got the power converter. He was able to find another of the proper voltage—it had belonged to a laptop that died in the heat—and used that until a power surge got it. Then one of his crew kludged up yet another power supply, which lasted until they were on the way home, before it too died. Note that the Z-Noteflex itself continued to work just fine through all this.

It's still working. The service department replaced the power supply under warranty, again no questions asked.

Phillip also notes that for the kind of use he put the machine to, it would have been better to have a bit thicker case, even at the cost of more weight; but barring a couple of chips and scratches, it was strong enough. Few machines are going to go through what that one did.

Meanwhile, I've stopped carrying my ancient Zenith Mastersport monochrome computer, which lasted me four years and went all over the world, including to the Mojave in midsummer. I wouldn't replace it with anything I didn't think equally reliable; but after Phillip's report, I trust the Z-Noteflex. The model I have is larger than Phillip's, and it has PCMCIA slots and an integral trackball. A Data Race RediCard RC-1496 data/fax modem takes up one slot. It works just fine.

One quibble: Zenith seems unable to make little doors that will stay on. The little cover on the PCMCIA slots keeps coming off, as did the cover over the serial and parallel ports on my Mastersport. On the Mastersport, I solved the problem by discarding the cover; the machine never missed it. The PCMCIA slots cover seems a bit more important, so I have applied the standard remedy: duct tape.

My Z-Noteflex has been on half a dozen trips by air and in my Bronco. It has not failed me. I could wish for a bit longer battery life, but I have to say it has never run out of steam on a cross-country flight; and in sleep mode (just close the lid), I can leave it for over a week before it quits.

This isn't the first User's Choice Award for Zenith laptops, and I doubt it will be the last. Zenith makes good PC/Windows portables.

Having said that, the best portable we have is Roberta's Mac PowerBook 540c, which is as powerful as the Quadra. It has great color

and sound, and it's just plain elegant. Of course, the only way I can use it is to pry it out of my wife's fingers; she really loves it. Apple's Mac PowerBook 540c wins a User's Choice Award and a Chaos Manor Orchid as well, and this isn't the last you'll hear of it.

The main product at Chaos Manor is words. I got my first computer (Zeke, who is now on display in the Smithsonian Museum of American History) to write books with Electric Pencil. I'm always looking for the perfect word processing software.

I haven't found it yet, so for the moment I use two different systems for writing. For first drafts of shorter works, I use Symantec's Q&A Write for DOS with Quinton Systems' Definitions/Plus, Soft-Key International's American Heritage Dictionary, and Microlytics' Word Finder Plus thesaurus, all loaded in a DOS session under Windows. It's what I'm writing this with now.

I use this for historical reasons: Q&A Write was the first really good text editor that incorporated the old WordStar control character system. This was a system in use back before keyboards routinely came with function keys. Many of us used the "WordStar Diamond" so often that we learned it down to the cellular level in our fingers.

Anyway, Q&A Write is a character-based, very fast text editor I find hard to improve on for first drafts, and the combination of the Q&A database and Q&A Write was one of the first "office suites" around. Q&A and Q&A Write are about the easiest to learn of the moderately powerful database and text-editing programs. I wrote lots of books and years of these columns with Q&A Write; but because, alas, Symantec has no intention of improving it, I can't in conscience recommend it any longer.

My other editor is Microsoft Word 6.0 for Windows, which I use for editing and printing. I often use it for creative writing, too. When I do that, I like it just fine, which naturally leads to the question: Why not use it all the time?

I'm used to Q&A Write. I don't often need a dictionary or thesaurus, but when I do, the combination works smoothly and how I want. I also plead efficiency. Somehow it seems like wretched excess to fire up a multimegabyte program like Word 6.0 to dash off a three-line letter. I can use Q&A Write on a DOS notebook or palmtop. Mostly, though, I have to confess it's

sheer inertia that keeps me from changing over.

The important thing is that if I had to give up one or the other, I'd give up Q&A Write. Word 6.0 has features I need for productivity. It handles footnotes and end notes. It compares different versions of a document. It does a lot of fancy formatting, allowing me considerable control over layouts. It does faxes, which I can send through WinFax directly from Word. I use Word a lot, and when I do first drafts with it, I like the fact that it changes *teh* to *the* (and a whole slew of other trivial error corrections) on the fly.

I had some problems with version 6.0 when it first came out. Actually I had a lot of problems. There were times when I was ready to tear out my hair. Now, though, version 6.0c is quite stable and reliable.

Moreover, Word integrates nicely with Excel and Access. Microsoft has done a nice job of interface design, and the Word/Access combination is as easy to use as the old Q&A/Q&A Write combination, only more powerful. I don't use Excel a lot, once again due to sheer inertia. I got used to SuperCalc when it first came out, and when I need a database, I've managed with Q&A, even though it's only a flat-file manager. Next year, however, I intend to reorganize, and Microsoft Office looks like the right way to go.

The User's Choice Award for Word Processor of the Year goes to Microsoft Word 6.0 for Windows.

Not to crow too loud, but I was about the first writer to recognize the value of CD-ROMs. Of course, the real credit should go to Bill Gates, who not only recognized their importance, but did something about it by hosting the first big CD-ROM conference.

CD-ROMs led to multimedia, and there are a lot of multimedia applications now. Many are games, but there are also lots of multimedia data disks, educational programs, maps, and so forth. There was a time when you could make do with a system with no CD-ROM and sound capabilities, but that time has passed.

Most new systems come with CD-ROM and sound capabilities, but a lot of older ones, some quite powerful, don't have them. As an example, Valiant had neither. Valiant is primarily an OS/2 machine and for much of this year ran OS/2 2.1. A few weeks ago I converted him to OS/2 Warp.

It was simple enough. Valiant already ran OS/2 LAN Server through an Eagle NE2000 Ethernet card, so I used the network to save everything onto a Maxoptix T3-1300 1.3-GB read/write optical disk drive, which is built into Pentafluge. It has always worked flawlessly, so much so that saving to the T3-1300 before trying anything radical has become standard procedure. The Maxoptix drive hangs off the SmartCache III SCSI Host Adapter. It works swiftly and flawlessly, and we don't have to think about it. It certainly deserves a User's Choice Award for Optical Disk Drives.

Once we had all the files saved, I reformatted Valiant's hard disk and installed PC-DOS 6 and the latest version of Windows for Workgroups. A caution here for ValuePoint owners. If you have a ValuePoint Pentium P5/60 with ATI Technologies' mach32 AX graphics system built onto the motherboard, you *must* put the statement ACEPLANAR=ON into your WIN .INI file in the section called mach. If you don't, you will never get out of VGA.

Second caution: if you have a network

card and W4WG, either be sure your Ethernet connection is working properly or bring up Windows with the /n option to prevent it from trying to start the network. In my case, I had disconnected the Ethernet cable from the back of the machine and later restarted and did "win." The result was a black screen of death, remedied only by reattaching the network cable. It may be that it will eventually reach a time out, but if so, that's a *long* timer.

Once I had W4WG running properly (including getting the machine connected to the network), I installed OS/2 Warp. Because the machine had no CD-ROM drive, I had to do it from floppy disks. This is tedious but uncomplicated, with one exception. If you have a network card in the system, once you get Warp going, you must use System Setup and do a Selective Install to get the adapter card set up. To do that, you have to tell OS/2 Warp what network card you have.

That turned out to be easier to say than do. An Eagle NE2000 card has absolutely no printing on it. If you don't know what kind of card it is, you can't find out by looking at it. Warp doesn't seem to recognize it automatically either.

Fortunately, the dual-boot capability of OS/2 Warp makes it easy to go back to DOS/Windows; and W4WG has, deep within the Windows Setup utility, a way to determine what network card you have. Open Windows Setup and invoke Change Network Settings. From there, open Drivers. In there, open Add Adapter and invoke Detect. The system will trundle for a while and tell you what network card it thinks it sees. W4WG will also set your network card for you (provided you have one that is software-settable; some are hard-wired with jumpers).

If you're installing OS/2 Warp, you will want to write down those settings. You can't use W4WG under Warp—not yet anyway—but if you leave the network card in, you still must have it set up right, or you cannot run Windows under Warp.

Another caution about W4WG. When I say it will set your network card for you, I mean it will do it whether you like it or not. When W4WG starts up, it looks for a network card set the way W4WG believes it to be set; if it finds a card with what it thinks are the "wrong" settings, it simply reprograms the card—and it does *not* tell you about that. You can get some spectacular crashes if W4WG believes your card should be set to some IRQ (interrupt request) or memory address used by another program.

Once I had the network card set, in-

stalling Warp was a snap; but I had no network capabilities, other than dual-booting back to DOS/Windows and using the still-working W4WG to transfer files. What I needed was an OS/2 LAN.

Fortunately I had one. Unfortunately, it was on CD-ROM. There was nothing for it. If I wanted to install OS/2 LAN Server, I'd have to put in a CD-ROM drive first.

I decided to try LANtastic for OS/2 first. I had some installation problems, which I can sum up in one sentence. Installing LANtastic on a DOS or Windows system is quite easy with the quick installation; but if you try to do a quick installation as opposed to a custom installation on an OS/2 Warp system, you will regret it. Use custom and pay attention. Also, remember the values your network card is set to.

Given that, LANtastic for OS/2 works. It doesn't work very well. It's fast enough once it establishes contact with a distant machine, but if it looks for and cannot find another machine on the network, it seems to take an eternity for it to recover, although it eventually does.

So long as you're running LANtastic for OS/2 from the OS/2 machine, you'll have no trouble transferring files to and from the machine. However, that's all you can do. If you go to a W4WG machine and try to connect to Valiant (or whatever you have named the OS/2 machine), it won't be able to find it. You can go back to Valiant and send—and receive—files from your W4WG machines, but that's it. If I wanted a full peer-to-peer network, I'd have to install LANtastic in my W4WG machines; otherwise, Valiant was a client but not a server.

That was a lot better than nothing, but it wasn't really convenient.

What I needed was a CD-ROM drive for OS/2. Fortunately I had one. Creative Labs recently sent me a Sound Blaster Multimedia Upgrade Kit, consisting of a double-speed CD-ROM drive, a Sound Blaster 16 card, and a pair of Labtec-powered speakers. It also comes with a lot of software, some of it rather old, but a lot of it excellent.

Installation was astonishingly easy. The drive mounted in Valiant's cabinet without a hitch. I had a problem with the cable. The installation diagram is designed for a different kit, I think, because it induced me to connect the data cable in reverse, with the result that the CD-ROM drive didn't work at all (and no wonder). Once I fixed that, everything went swimmingly, and I now have good CD-ROM and sound capabilities; in a word, Warped Valiant has become a multimedia system.

The sound picture remains confused, and I keep hearing about new developments coming Real Soon Now. That's fine, but you need CD-ROM and sound capabilities right now. The way to get them is with one of Creative's excellent Sound Blaster kits, which work under DOS/Windows even better than under OS/2 Warp. Creative Labs' Sound Blaster Kits get the User's Choice Award for Multimedia Upgrade Kits.

Once I had the CD-ROM drive installed, I put in OS/2 LAN Server. A long story goes with that, too, but I don't have room for it. The short version is, *do not* use the express or quick installation. Use the custom installation and be prepared for some confusing screens. The express method believes your network board is set to IRQ 3, which isn't a good setting. It offers you no way to tell the program where the board is really set.

If you give LAN Server the name of a W4WG workgroup—mine is JerryOne— as the domain, it's very easy to make the OS/2 computer part of a W4WG network; I can find Valiant from any machine in the network. It's a fast and smooth process. Incidentally, if you want to defeat the network security, give the guest account system administrator privileges. You do that in the network settings book.

I am out of space, and I haven't got to the Orchids and Onions. For that matter, some of the awards still need to be passed out, so some of this will continue next month.

Two games get User's Choice Awards: MicroProse's XCOM:UFO Defense, which is the best strategy game I've played in a long time, and Origin's Wing Commander 3. WC 3 is awesome, with lots of live-action scenes involving believable characters and a good story line. One warning: there are some missions I just can't manage. I'm told you can complete the game on the winning track even though you lose those missions.

Logitech gets a User's Choice Award for its WingMan Extreme joystick. It's the best one I've come across, and I've used a lot of them.

The book of the month is *Technological Risk* by H. W. Lewis (Norton, 1992). It's about the clearest and best-written exposition on the increasingly important subject of the risks involved with new technologies. It's readable, too. The computer book of the month is by Gary Wolf and Michael Stein, *Aether Madness* (Peachpit Press, 1994). This is a wacky guide to places to go and things to look at on the Internet. The glossary has readable explanations of `ftp`, archie, veronica, gopher, and other esoterica, and the main body of the book will tell you about odd stuff indeed.

Next month, Orchids and Onions, a few more awards, and what's new at Chaos Manor. ∎

Jerry Pournelle holds a doctorate in psychology and is a science fiction writer who also earns a comfortable living writing about computers present and future. Jerry welcomes readers' comments and opinions. Send a self-addressed, stamped envelope to Jerry Pournelle, c/o BYTE, One Phoenix Mill Lane, Peterborough, NH 03458. Please put your address on the letter as well as on the envelope. Due to the high volume of letters, Jerry cannot guarantee a personal reply. You can also contact him on the Internet or BIX at jerryp@bix.com.

For More Information

I can't live without **BOOTCON 2.11** ($49), and if you play games or experiment with DOS configurations, you can't either. Contact **Modular Software Systems**, Kent, WA, (800) 438-3930 or (206) 631-5781; fax (206) 631-5779. **Circle 1146 on Inquiry Card.**

Crescent Software offers a bunch of programming tools, such as **Graphics QuickScreen 1.12** ($149), **QuickPak Professional 4.19** ($199), and **QuickPak Professional 3.22 for Windows** ($199), and I'm giving a User's Choice Award in Programming Utilities to the whole line. Contact **Crescent Software, Inc.**, Ridgefield, CT, (800) 352-2742 or (203) 438-5300; fax (203) 431-4626. **Circle 1147.**

So long as you're running **LANtastic for OS/2** (per node, $139; five-user pack, $629; 10-user pack, $1119; 25-user pack, $1919; 50-user pack, $3129; 100-user pack, $4995) from an OS/2 machine, you'll have no trouble transferring files to and from the machine. Contact **Artisoft, Inc.**, Tucson, AZ, (800) 233-5564 or (602) 670-7100; fax (602) 670-7101. **Circle 1148.**

The **Mac PowerBook 540c** (base price, $4299) has great color and sound, and it's just plain elegant. Contact **Apple Computer, Inc.**, Cupertino, CA, (800) 776-2333 or (408) 996-1010; fax (904) 584-7481. **Circle 1149.**

Microsoft Word 6.0 for Windows ($339) handles footnotes and endnotes, compares different versions of a document, and does a lot of fancy formatting and faxes. Contact **Microsoft Corp.**, Redmond, WA, (800) 429-9400 or (206) 882-8080; fax (206) 883-8101. **Circle 1150.**

OS/2 Warp 3 for Windows (less than $80) and **OS/2 Warp 3 Full Pack Edition** (less than $130) aren't perfect, but for many of us they're more than good enough. Contact **IBM Corp.**, Austin, TX, (800) 342-6672 or call your local IBM dealer. **Circle 1151.**

I'd be lost without **QEMM 7.5** ($99.95), and **SideBar** ($59.95) was my favorite Windows utility of the year. Contact **Quarterdeck Office Systems**, Santa Monica, CA, (800) 354-3222 or (310) 392-9851; fax (310) 314-4217. **Circle 1152.**

The way to get CD-ROM and sound capabilities right now is with a **Sound Blaster Multimedia Upgrade Kit** ($340 to $599). Contact **Creative Labs, Inc.**, Milpitas, CA, (800) 998-1000 or (408) 428-6600; fax (408) 428-6611. **Circle 1153.**

The 1.3-GB, multifunction **T3-1300** optical disk drive (with 4 MB of cache memory: internal, $2999; external, $3285) works swiftly and flawlessly, and we never have to think about it. Contact **Maxoptix Corp.**, San Jose, CA, (800) 848-3092 or (408) 954-9700; fax (408) 954-9711. **Circle 1154.**

The most powerful machine at Chaos Manor is Pentafluge, a Pentium system that includes a PC Power & Cooling **tower case** ($279), **Turbo-Cool 300 tower-model power supply** ($179), and **Pentacool fan** ($29). Contact **PC Power & Cooling, Inc.**, Carlsbad, CA, (800) 722-6555 or (619) 931-5700; fax (619) 931-6988. **Circle 1155.**

Vopt 4.02 ($59.95) has always been the safest disk defragger I know of. Contact **Golden Bow Systems, Inc.**, San Diego, CA, (800) 284-3269 or (619) 298-9349; fax (619) 298-9950. **Circle 1156.**

Wing Commander 3 ($65 to $75) is awesome, with lots of live-action scenes involving believable characters and a good story line. Contact **Origin Systems, Inc.**, Austin, TX, (800) 245-4525 or (512) 335-5200; fax (512) 331-9558. **Circle 1157.**

The **WingMan Extreme** ($69.95) joystick is the best one I've come across, and I've used a lot of them. Contact **Logitech, Inc.**, Fremont, CA, (800) 231-7717 or (510) 795-8500; fax (510) 792-8901. **Circle 1158.**

XCOM: UFO Defense (price not available) is the best strategy game I've played in a long time. Contact **MicroProse Software**, Hunt Valley, MD, (800) 879-7529 or (410) 771-1151; fax (410) 771-1174. **Circle 1159.**

My **Z-Noteflex** (8 MB of RAM, dual-scan screen, PCMCIA slots, and an integral trackball, $3379) has been on half a dozen trips by air and in my Bronco. It has never failed me. Contact **Zenith Data Systems Corp.**, Buffalo Grove, IL, (800) 533-0331 or (708) 808-5000; fax (708) 808-4482. **Circle 1160.**

Circle 70 on Inquiry Card.

HEAVY ON POWER

Weighing 8½ pounds, the base-model PowerLite 85 notebook offers 85-MHz MicroSparc II performance, 510 MB of disk storage (expandable to 4.4 GB), and a TGX-accelerated 640- by 480-pixel or 1024- by 760-pixel display. Among its other features are a full-size keyboard, 32 MB of memory (expandable to 96 MB), an internal floppy drive and fax modem, and Solaris 1.1.2. Pricing starts at $11,995.

Contact: RDI Computer, Carlsbad, CA, (800) 734-5483 or (619) 929-0992.

Circle 1271 on Inquiry Card.

MIX-AND-MATCH STORAGE

The Mega Drive Enterprise, from Mega Drive Systems (Beverly Hills, CA), is an intelligent storage system that has the ability to handle hard drives, optical storage, CD-ROMs, removable tapes, and DAT. Fiber-ready and compatible with most SCSI host systems, the product includes two independent channels that offer 40-MBps throughput using Fast and Wide SCSI-2. With a fiber-channel interface, the system's throughput reaches 200 MBps. The eight-slot system ($1200, unloaded) is available in tower or rack-mount configurations.
Phone: (800) 404-6342 or (310) 247-0006.
Circle 1275 on Inquiry Card.

FOUR-WAY VIDEOCONFERENCING

With the C-Phone Multipoint Controller ($3995), from Target Technologies (Wilmington, NC), you can use the C-Phone video PBX system ($1995) to view four other people on your screen during a videoconference call or display an alternative video source in one or more of the four windows. Participants can be connected to your network or call in from remote locations. You can either shrink the full-screen display to a portion of the Windows desktop or enlarge one image to full-screen size. Audio for all participants is synchronous.
Phone: (800) 666-2496 or (910) 395-6100.
Circle 1274 on Inquiry Card.

ETHERNET CARDS FOR THE MAC

Dayna Communications (Salt Lake City, UT) offers two DaynaPort E/CS Ethernet NICs for Macs equipped with a communications slot. The 10Base-T version, the E/CS-T, costs $109; the thin Ethernet version, the E/CS, costs $119. Both cards are Apple Driver–compatible.
Phone: (801) 269-7200.
Circle 1286 on Inquiry Card.

UNIVERSAL MODEMS ▲

Unlimited Systems (San Diego, CA) offers the Konexx Konnector Model 111 ($159) and the Konexx Koupler Model 204 ($149). Both provide high-speed data connections to digital PBX systems and most hotel, cellular, and international phones. To use the 28.8-Kbps Model 111, you unplug the coiled handset cord from the telephone base unit, plug in the Model 111, and plug the handset cord into the Model 111. The 24-Kbps Model 204 is an acoustic coupler that you plug into a telephone-line jack on a modem and strap onto a telephone handset.
Phone: (800) 275-6354 or (619) 622-1400.
Circle 1278 on Inquiry Card.

MINI-CARTRIDGE BACKUP

A 3½-inch tape backup device from Tandberg Data (Simi Valley, CA), the Panther Mini 1000 (from $799) comes in both internal and external models. The complete subsystem features a SCSI-2 tape drive, a data cartridge, cables, a user's manual, and software for DOS/Novell/Windows, OS/2, or Unix/Windows NT. Using a Sony wide-tape cartridge, you can increase the Panther Mini 1000 system's standard 1-GB storage capacity to 2 GB. Its backup speeds can reach 18 MB per minute, and a read-after-write feature ensures data integrity.
Phone: (800) 826-3237 or (805) 579-1000.
Circle 1276 on Inquiry Card.

SERVING UP PRINTING

The family of multiprotocol Token Ring printer servers from DPI (Waltham, MA) work with NetWare and Vines networks. The NetPrint/50T ($595) supports printing from NetWare and Unix TCP/IP to a parallel printer, and the NetPrint/100T ($795) can print to parallel and serial printers. The NetPrint/200T for NetWare ($1295) also supports TokenTalk and permits printing to one serial and three parallel printers. The Vines printer servers—the NetPrint/104T ($1095), the NetPrint/200T for Vines ($1495), and the JetxPrint/104T ($1195)—allow users of Vines, Unix, and NetWare to concurrently print to the same printers using their native operating systems.
Phone: (800) 243-2333 or (617) 647-1234.
Circle 1289 on Inquiry Card.

COMBINED POWER

Designed for Mach DS plus LAN multiplexers, the Mach Terminal Server module ($2495), from Data Race (San Antonio, TX), combines the functions of a statistical multiplexer with an IP card and a LAT terminal server. It lets local and remote terminals access host systems on a single Ethernet connection. Other advantages include service to multiple remote sites with one multiplexer; dynamic allocation of bandwidth to let data, voice, and LAN traffic share a composite link; and an SNMP-compatible interface.
Phone: (800) 329-7223 or (210) 558-1929.
Circle 1279 on Inquiry Card.

MODEMLESS COMMUNICATIONS

Telebyte Technology's (Greenlawn, NY) Model 234 is an RS-232 modem eliminator designed to be placed between two asynchronous or synchronous DTEs. This permits them to communicate over a limited distance without a pair of synchronous modems. The unit ($185) has two DB-25S female connectors, supports nine RS-232 leads, and operates at speeds from 2.4 Kbps to 38.4 Kbps. It draws its power from the data and control leads on the DTE connectors.
Phone: (800) 835-3298 or (516) 423-3232.
Circle 1281 on Inquiry Card. ▼

BRIGHTER IMAGES, CLEARER ▲ SOUNDS

The Bedini Clarifier ($39.95), from Bedini Electronics (Coeur d'Alene, ID), provides a low-cost solution to removing unwanted electrical charges from the surface of CDs. The 1-pound device polarizes stored data, maximizing color resolution and brightness and enhancing sound.
Phone: (800) 876-0299 or (208) 667-8300.
Circle 1285 on Inquiry Card.

ULTRALIGHT VIRTUAL REALITY ▲

Virtual i-O (Seattle, WA) offers its Virtual i-O i-glasses, which are lightweight, high-tech 3-D glasses that create the illusion of a large virtual screen floating in front of you. The company offers an 8-ounce base model ($599) for TV and video-gaming and an advanced model ($799) that uses head-tracking technology to provide a 360-degree environment and a 30-degree horizontal field of view.
Phone: (800) 646-3759 or (206) 382-7410.
Circle 1293 on Inquiry Card.

IMPROVING REMOTE COMMUNICATIONS

An intelligent serial-port board for remote access, the AccelePort, from DigiBoard (Eden Prairie, MN), provides 230-Kbps unidirectional communications, 115-Kbps bidirectional communica-

tions, and a single point through which to control traffic. The product features on-board serial processing and lets you connect as many as eight asynchronous modems to an ISA file server. A four-port AccelePort costs $895; an eight-port model is $1095.
Phone: (800) 344-4273 or (612) 943-9020.
Circle 1308 on Inquiry Card.

HIGH-OCTANE DATA TRANSFERS

The Multiport Model 6400SX, from Aurora Technologies (Waltham, MA), enables users of SPARC workstations and servers to add 64 high-speed asynchronous serial lines to their networks. The device (from $5995) provides higher throughput and lower system overhead by processing I/O on-board. It links modems, DSUs, and other RS-232, RS-422, or RS-485 serial devices, providing transfer rates as high as 115.2 Kbps from a single host SBus slot.
Phone: (617) 290-4800.
Circle 1283 on Inquiry Card.

A SMALL, RUGGED FLASH DISK

From M-Systems (Santa Clara, CA), the PC/104 Flash Disk provides 1 to 32 MB of storage in a 3.6- by 3.8-inch package. The disk ($99.95), which plugs into a standard PC/104 connector, operates reliably under conditions such as severe shock, humidity, dust, and extreme temperatures. The unit offers performance fast-

er than that of a rotating disk but consumes less power. In addition, fragmentation doesn't degrade its performance. The on-board BIOS eliminates the need for installing software.
Phone: (408) 654-5820.
Circle 1284 on Inquiry Card.

LOW-COST MONITORS

Sampo Technology (Norcross, GA) offers a line of corporate-grade monitors. The 14-inch noninterlaced AlphaScan SV KM-400 ($269) has a dot pitch of .28 mm, is Energy Star–compliant, works with Macs and PCs, and offers resolutions as high as 1024 by 768 pixels for VGA and 832 by 624 pixels on the Macintosh. The 15-inch AlphaScan 15mx KM-510 ($389) and the 17-inch Alpha-Scan 17mx KM-710 ($699) have similar characteristics, but they permit a maximum resolution of 1280 by 1024 pixels for VGA and 1024 by 768 pixels on the Macintosh.
Phone: (404) 449-6220.
Circle 1290 on Inquiry Card.

LINK LAPTOPS TO DESKTOPS

An infrared device, the JetEye PC ESI-9680A, from Extended Systems (Boise, ID), lets you transfer data from a desktop to a laptop without using a cable. The unit works with any portable that supports the IrDA infrared standard. You attach the Jet-

Eye ($135) transceiver to the serial port on your desktop. To establish a connection, you aim your laptop's infrared port at the transceiver. You can access data and applications from a floppy disk or hard drive or send data to a printer. The Jet-Eye operates at 115.2 Kbps.
Phone: (800) 235-7576 or (208) 322-7575.
Circle 1288 on Inquiry Card.

FAST, TRUE-COLOR ▲ GRAPHICS ACCELERATION

A 24-bit true-color graphics accelerator card for PCs, the Max-media CX/Pro ($999) uses three 64-bit processors to give speedy performance. From UMax Technologies (Fremont, CA), the card supports a maximum resolution of 1600 by 1200 pixels, can display true color when operating at refresh rates as high as 120 Hz, and uses inexpensive DRAM chips. Running Windows and Autodesk, the Maxmedia card enlarges the viewing area beyond the screen's limits, creating a virtual-desktop area of up to 2048 by 1024 pixels.
Phone: (800) 562-0311 or (510) 651-8883.
Circle 1282 on Inquiry Card.

BROADCAST PRINTING

A wireless printer sharer that's compatible with DOS and Windows, GoPrint transmits data at speeds as high as 1 Mbps. Using spread-spectrum radio, the unit can broadcast a signal as far as 3000 feet through nonobstructive material. To set up the unit, you plug its transceivers ($399 per pair) into the parallel ports of a computer and printer. Several computers can share a printer, and several printers using GoPrint can share an office.
Contact: AeroComm, Lenexa, KS, (800) 492-2320 or (913) 492-2320.
Circle 1272 on Inquiry Card.

RECORDABLE CD FOR UNDER $2000

The RCD-1000 is a 5¼-inch Macintosh recordable CD system that offers double-speed reads and writes. Equipped with a 1-MB cache buffer, the unit supports ISO 9660, HFS, digital audio/video, multisession, and mixed-mode formats. Each CD offers random access and holds as much as 650 MB of data or 74 minutes of audio. The $1995 package includes authoring software and a backup utility.

Contact: Pinnacle Micro, Irvine, CA, (800) 553-7070 or (714) 789-3000.

Circle 1273 on Inquiry Card.

FAULT PREVENTION

Fiber•Mau-FDX ($299), a full-duplex Ethernet transceiver from MiLAN Technology (Sunnyvale, CA), uses the company's Link-Sentry remote-sensing technology to search for errors on a network before releasing packets to their destinations. LinkSentry handles this task without disturbing network traffic. If any cabling or equipment problems arise, the fault-tolerant transceiver or hub is notified to disable the transmitting port. Networks running SNMP can detect the disabled port.
Phone: (408) 752-2770.
Circle 1280 on Inquiry Card.

LONG-DISTANCE COLOR PRINTING ▼

The ColorBlaster, from PRDS (Fremont, CA), handles color PostScript files sent over phone lines via a modem. This allows Mac, PC, and workstation users to print their color documents, graphics images, and charts from miles away on expensive color printers or on color copiers that are equipped with EFI's Fiery.

The unit ($1495) comes with a 14.4-Kbps modem and BBS software that requests identification and billing information from a caller before permitting the caller to send files. ColorBlaster automatically processes PostScript files and converts fax files to PostScript format.
Phone: (510) 226-8930.
Circle 1292 on Inquiry Card.

STORAGE FOR PCS AND MACS

The Advantage series of storage products, from MicroNet Technology (Irvine, CA), now supports PC compatibles running DOS or Windows, as well as Macs and Power Macs. To complement the line's PC compatibility, Micro-Net is marketing the high-performance Advantage Host Adapter ($99) for ISA systems, which comes with Tape Mate II backup software, as well as software for connecting SCSI devices. Prices for the Advantage hard drives start at $329 for a 170-MB unit; prices for a SyQuest-based removable drive start at $499 for 88-MB capacity.
Phone: (714) 453-6100.
Circle 1309 on Inquiry Card.

HIGH-CAPACITY TAPE BACKUP

The Legacy QIC 10 drive, from Legacy Storage Systems (Andover, MA), is a 5¼-inch, half-height tape unit capable of storing 10 GB of compressed data on a single quarter-inch cartridge. Hardware data compression enables the unit to back up 1 GB of data in 11 minutes; the SCSI-2 device can locate files in an average of 40 seconds. You can connect the external model ($1800) or the internal model ($2098) to a PC, workstation, or server. The internal device also fits in a Legacy SmartArray disk-array subsystem.
Phone: (508) 681-8400.
Circle 1294 on Inquiry Card.

COMMUNICATIONS ON THE DOUBLE

The V4-ISA-650T-RJ multiport adapter, from Computone (Roswell, GA), lets you add high-speed RS-232 serial ports to an ISA-based DOS or Unix system. Equipped with 16C650 UARTs, the adapter ($495) offers superior buffering and sustained full-duplex data transfer rates of 460 Kbps. The unit has four built-in RJ-45 connectors, complete surge/spike protection, and full modem control on each port. The four ports can share a single IRQ, and you can install two adapters in the same system.
Phone: (800) 241-3946 or (404) 475-2725.
Circle 1291 on Inquiry Card.

HIGH-SPEED SCSI-2 ▲ HOST ADAPTERS

BusLogic (Santa Clara, CA) offers two Fast and Wide SCSI-2 host adapters for PCI bus systems: the BT-956C ($399) and the BT-956CD ($599). Both support up to 15 SCSI devices; can handle Fast and Wide disk drives while running 8-bit devices, such as CD-ROMs; are PCI 2.0–compliant; and feature 32-bit bus-master DMA transfers as fast as 133 MBps across the PCI in burst mode. The BT-956C has built-in support for the emerging Ultra SCSI specification.
Phone: (408) 492-9090.
Circle 1277 on Inquiry Card.

CONSERVE ENERGY ▲

To reduce the amount of electricity that's consumed by computer monitors, Merritt Computer Products (Dallas, TX) offers PowerMan. The $99.95 package tracks the activity of your keyboard and mouse, turning the monitor off after a specified period of inactivity. You receive both a visual and an audio warning before shutdown. In addition, PowerMan provides password protection and improved data security and lets you turn off the monitor by pressing a key combination. The device works with VGA and SVGA monitors.
Phone: (800) 627-7752 or (214) 339-0753.
Circle 1287 on Inquiry Card.

FOR YOUR EYES ONLY

The wraparound design of the PF50 filter ($119), from 3M Optical Systems (Roseville, MN), preserves the privacy of travelers who use their notebook computers in public places. The screen is legible only when viewed from directly in front of the display. An antiglare feature improves contrast and reduces glare.
Phone: (800) 553-9215 or (612) 733-4403.
Circle 1310 on Inquiry Card.

World Dominance

Business Week (US)

Data Communications (US)

Software Digest (US)

PC Digest (US)

Personal Computer Magazine (UK)

BYTE (US)

Datatid (Norway)

L'Ordinateur individuel (France)

Personal Computer Magazine (Netherlands)

Datanews (Belgium)

BAJT (Czech-Slovak Republics)

Nikkei BYTE (Japan)

LAN Times (US)

BYTE Mexico

Open Computing (US)

BYTE Argentina

BYTE Brasil

Binary (Spain)

Bit (Italy)

RAM (Greece)

BYTE Bulgaria

BYTE Croatia

BYTE Türkiye

BYTE Middle East

0&1/BYTE (Taiwan)

NSTL Benchmarks Are The Worldwide Standard

From North America to Asia, South America to Europe, more than 20 publishing partners utilize test results produced by NSTL's unique, cross platform InterMark™ benchmarks, providing more than 8 million readers globally with objective, real world product evaluations.

Think global — and put NSTL to the test.

> Look for NSTL's InterMark test results in **Business Week**, **BYTE**, **Data Communications**, **LAN Times**, **Open Computing**, **PC Digest** and **Software Digest**. For more information call 1-610-941-9600.

MEETS TEST STANDARDS

NSTL

★★★★★

NSTL • 625 Ridge Pike • Conshohocken, PA 19428 • 1- 610 - 941-9600 • Fax: 1-610 -941-9952 • NSTL is a division of McGraw-Hill, Inc.

What's New Software

GET ORGANIZED

DragStrip lets you arrange Mac programs and documents on draggable strips and then launch them with a single mouse-click. Applications store the most recently launched documents in a pop-up menu off of each application on your strips. Dragging a folder or a hard drive onto a strip displays a hierarchical pop-up menu; you click on an item to view it. DragStrip supports Apple ControlStrip modules and includes its own add-ons, called DragStrip Additions, for such operations as changing the system's sound volume or the monitor's bit depth. Accelerated for Power Macs, the program costs $59.95.

Contact: Natural Intelligence, Cambridge, MA, (800) 999-4649 or (617) 876-4876.

Circle 1296 on Inquiry Card.

THE MORE DATA, THE MERRIER

capaCD Pro, from EWB (Carlsbad, CA), raises the 660-MB storage limit on CD-ROM capacity, letting you record, on average, two to three times more information per disc. After analyzing the types of data to be stored, capaCD Pro uses the optimal compression technique for each type. The program (from $99.95 for a 100-platter distribution license) uses industry-standard media; the enhanced discs can also be used in ordinary CD-ROM drives.
Phone: (619) 930-0440.
Circle 1299 on Inquiry Card.

PROFESSIONAL CAD

Mental Automation's (Bellevue, WA) SuperCAD+ for Windows, a schematic editor, offers pull-down menus, dialog boxes, and toolbars. Compatible with MDI, the program ($249) lets you load as many as 20 schematic sheets simultaneously. Editing features include automatic bus generation, snap to grid, automatic wiring,
and title-block annotation. You can add your own symbols to the library's 2600 schematic symbols, change the drawing scale, and zoom in on a particular area.
Phone: (206) 641-2141.
Circle 1311 on Inquiry Card.

ALL BUTTONED UP ▲

MDA (Cincinnati, OH) offers two sets of electronic buttons, from its Electronic Preview Collection, for use with any multimedia authoring program or presentation package. The animated buttons are spherical glass objects resembling marbles with unusual textures; they come in pushed and unpushed styles. The neon buttons resemble miniature neon signs and come in "on" and "off" styles. For PCs and Macs, the files are offered as 24- or 8-bit TIF and PICT files. Each 50-
button set costs $59; there's no royalty fee.
Phone: (800) 359-2964 or (812) 579-5063.
Circle 1301 on Inquiry Card.

BLANKET PROTECTION

Developed by Atemi (Champaign, IL) for the Mac, Night-Shade (from $199) encrypts all communications on LocalTalk, Ethernet, and Token Ring networks. You set your preferences, and encryption takes place transparently. The program works with ARA 1.0 and 2.0, as well as across multiple zones and on internets. The program allows unencrypted sessions with machines not running NightShade. A DES or triple-DES module costs $99; an AppleShare NLM version is scheduled for release this month.
Phone: (800) 542-8364 or (217) 352-3688.
Circle 1302 on Inquiry Card.

GROUP INTERACTION

Built on the OLE 2.0 distributed container model, Insitu Conference (from $269), from Insitu (Boston, MA), provides a means for groups to generate ideas, edit documents, and discuss an image or document displayed on a Windows desktop. Its features include multiparticipant document editing and file transfer, automated telephony conferencing, and a shared whiteboard, on which each group member can write using one of 16 pens of different colors.
Phone: (617) 720-0821.
Circle 1306 on Inquiry Card.

FASTER AUTOCAD

From Vermont Microsystems (Winooski, VT), AutoMate/Pro speeds up the graphics in releases 12 and 13 of Autodesk's Auto-CAD for Windows. The product improves AutoCAD release 13's pans, zooms, and redraws by as much as 10 times and betters the
speed of erasing and moves by as much as 300 percent. In addition, AutoMate/Pro (from $295) offers a graphical interface for Auto-CAD 12 and 13 for DOS.
Phone: (800) 354-0055 or (802) 655-2860.
Circle 1300 on Inquiry Card.

IN LIVING COLOR ▲

Color Compass ($129) from Praxisoft (Ashburn, VA) simplifies the task of creating, modifying, and blending colors on the Mac. Using the software's two-, three-, and four-color blending windows, you can create and manipulate colors by changing the RGB or CMYK mix. You can also drag between windows and palettes and maintain links to permit instantaneous updates. Other features include Superselect mode, which lets the palettes control the blends; the ability to display color families; and blend reweighting. The package is accelerated for Power Macs.
Phone: (800) 557-7294 or (703) 729-3391.
Circle 1304 on Inquiry Card.

A NEW TWIST ON LOTUS

Multi-D Lite ($249), a workstation version of the RServer multidimensional database server and RGUI graphical shell, makes it easier to extract information from a database and drop it into Lotus 1-2-3 release 4 or Lotus 1-2-3 for OS/2 spreadsheets. From Renaissance Information Technology (Grapevine, TX), the native 32-bit OS/2 application lets you filter data and display it in a variety of ways.
Phone: (817) 421-8127.
Circle 1305 on Inquiry Card.

NETWORK ASSISTANT

More than a troubleshooting tool, NetReport (from $495) provides you with information critical to long-term capacity planning. From HawkNet (Carlsbad, CA), the program tracks over 110 NetWare internals, providing access to real-time data and compiling historical data for up to 90 days. Customized Windows screens display details on every file server on the network. To analyze network activity, you select from 12 reporting categories, including licensing levels and revision numbers for NetWare, memory, disks, and NLMs.
Phone: (800) 429-5638 or (619) 929-9966.
Circle 1312 on Inquiry Card.

STUCK ON YOU

StickAround Notes ($24.95) from O'Soft (Cambridge, MA) lets you put electronic notes on your Windows display. The heart of the program is a note dispenser that's always on the screen. A single mouse-click makes a blank note pop up; everything you type on a note is stored automatically. Each time you start Windows, your notes appear exactly where you last left them.
Phone: (800) 888-4437 or (617) 225-2010.
Circle 1317 on Inquiry Card.

REAL-TIME ▲ PROGRAMMING TOOLS

From Quinn-Curtis (Needham, MA), Real-Time Graphics Tools for Windows lets you create real-time displays in applications running under the Windows family of operating systems. The kit includes dynamic graphs that you update in real time with smooth scrolling and only a minimum of redrawing and flicker. Dynamic graph types include scrolling, sweep, bar, and logic graphs; *x,y* plots; and meters. The base versions, for C/C++ ($600) and Visual Basic ($500), offer 16- and 32-bit DLLs. The advanced developer's versions, for C/C++ ($1200) and Visual Basic ($1100), contain source code.
Phone: (617) 449-6155.
Circle 1315 on Inquiry Card.

PLANNING AHEAD

PlanMaker, from PowerSolutions for Business (St. Louis, MO), helps you prepare, formulate, and publish a marketable business plan. It prompts you to answer questions, shows you how to present concepts, and asks for financial data. From there, it can produce such reports as a 12-month pro forma invoice, a three-year projected operating statement, a cash-flow projection, and a balance sheet. Available in DOS and Mac versions, PlanMaker ($129) comes with an on-screen tutorial.
Phone: (800) 955-3337 or (312) 421-0670.
Circle 1307 on Inquiry Card.

TUNING IN, GETTING ON

ICL (Reston, VA) offers Embla, a Windows application that provides access to Internet mail for PC users on Unix networks. In addition to the standard collection of mail functions for documents, the product (from $99) offers support for sending and receiving spreadsheets, multimedia images, and sound files. Embla automatically alerts you to incoming mail and stores messages in folders located either on your PC or on a Unix server. The package includes an editor and enables you to create aliases and distribution groups.
Phone: (703) 648-3300.
Circle 1303 on Inquiry Card.

LOW-COST UNIX CD-ROM

A Unix operating system for 386 and 486 PCs, Linux includes complete source code, programming languages, hundreds of megabytes of utilities, and a printed installation manual. Developed on the Internet, the $69 CD-ROM package comes with drivers for sound cards, CD-ROMs, video cards, IDE interfaces, and SCSI host adapters. It also

includes a demo version of FlagShip, a database development system and one of the first native commercial Linux applications. Linux provides support for TCP/IP, NFS, DOS emulation, FTP, SLIP, PPP, and Telnet.
Contact: WorkGroup Solutions, Aurora, CO, (800) 234-7813 or (303) 699-7470.
Circle 1295 on Inquiry Card.

In the Changing World of Technology, One Size Does Not Fit All

Made in the U.S.A.

From Notebook to Desktop

Quatech's PCD2-F Internal Interface Adapter allows you to add PCMCIA capability to your desktop. The adapter is easy to install, and supports Type I, Type II and Type III memory and I/O cards. The PCD2-F and Quatech's PCMCIA I/O cards conform to 2.1 Specifications. The I/O cards support hot-swapping and include configuration software.

Communication cards for PCMCIA include single channel and dual channel RS-232 or RS-422/485 asynchronous serial adapters, a multi-protocol adapter, and an enhanced parallel port. Check out Quatech's complete line of communication and data acquisition cards for PCMCIA.

High-speed Communication Adapters

Quatech's communication adapters for the AT bus meet asynchronous or synchronous and serial or parallel communication requirements with protocols such as RS-232, RS-422, RS-485, Current Loop, and IEEE488.

Expand your system with the DS-100, a dual channel RS-232 adapter. The QS-100D is a four channel RS-232 adapter. Both feature optional 16550 UARTS, all AT interrupts and flexible addressing. The QS-100D supports OS/2, UNIX, Xenix, DOS, and Windows®. The MPA-102 is a dual channel multi-protocol adapter utilizing an AMD 85C30-10 SCC. Other communication adapters for the AT bus are also available.

QUATECH

Call **1-800-553-1170** for more information and a FREE handbook of our complete line of PCMCIA, communication and data acquisition products.

Foreign Distributor Inquiries Welcome

Made in the U.S.A. 662 Wolf Ledges Parkway, Akron, Ohio 44311 (Phone 216-434-3154/Fax 216-434-1409). International Distributors: **Australia**/Interworld Electronics 61-3-563-5011, **Belgium**/ Acal NV/SA 32-27-205983, **Brazil**/Intercomp Electronica Ltda (Sao Paulo) 55-11-8532733, Medusa (Rio de Janeiro) 55-21-2554745, **Canada** (Western)/Interworld VCR 604-984-4171 (Toronto 905-513-7027), **Denmark**/Jes Rasmussen Aps. 45-4281-6838, **England**/Diamond Point International 44-634-722-390, **Finland**/Lab Hitech OY 358-0-682-1255, **France**/Elexo 33-1-69302880, **Germany**/Jupiter Electronic Systems 49-61-8175041, **Hong Kong**/Brio Technology Ltd. 852-581-1111, **India**/Computaccount 91-11-224-5159, **Israel**/RCM Ltd. 972-03-6487885, **Italy**/N.C.S. Computer Italia 39-331-770016, **Japan**/Nictrix Corporation (New Jersey) 201-947-2220, **Korea**/Sam Boo Enterprise Co. 82-2-5384001, **Netherlands**/ACAL Auriema 31-40-502602, **Saudi Arabia**/Integrated Computer Operations 966-3-895-1827, **Singapore**/Bliss Services Pte Ltd 65-338-1300, **South Africa**/Eagle Technology 27-21-234943, **Spain**/SANTA Barbara SA 34-3-418-81-16, **Sweden**/SYSTEC 46-13-1101-40, **Switzerland**/Technosoftware Ltd. 41-64-519040.

Circle 96 on Inquiry Card (RESELLERS: 97).

BYTE

BUYER'S GUIDE

Essential Products and Services for Technology Experts

Mail Order

Top mail-order vendors offer the latest hardware and software products at the best prices.

240

Hardware/Software Showcase

Your full-color guide to in-demand hardware and software products, categorized for quick access.

265

Buyer's Mart

The BYTE classified directory of computer products and services, organized by subject so you can easily locate the right product.

275

Up to 4 Users can Access Multiple PC, Macintosh and Sun Computers from a Central Location!

● Multiuser

Up to 4 users have simultaneous access to any attached computer!

● Multiplatform

Supports any combination of PC, Macintosh and Sun computers; use any platform's peripherals to access any type of computer in the system.

● Multimedia

Every user has full multimedia capabilities; supports keyboard, mouse, video, speakers, microphone and serial port.

COME SEE US AT
CeBIT in Hannover, Germany, March 8-15, 1995 Booth #M9, M10
Networld & Interop in Las Vegas NV, March 28-30, Booth # 5839
COMDEX Spring in Atlanta, GA, April 24-27 1995 Booth #N5762

AUTOBOOT COMMANDER 4xP™

Cybex Corporation
4912 Research Drive • Huntsville, AL 35805 USA
(205) 430-4000 • FAX (205) 430-4030

CYBEX™

Circle 166 on Inquiry Card (RESELLERS: 167).

PLUG

One Portable CD-ROM Drive — Two Simple Steps

PLAY

With the backpack quad-speed CD-ROM drive, you're one cable connection away from a world of multimedia business advantages: Plug it into your PC parallel port to access a universe of CD-ROM titles. Take it with you for portable data retrieval. Or share it with other PCs.

No compatibility problems with virtually any IBM-compatible PC— Windows, DOS or OS/2. No taking your PC apart to install this quad-speed drive. Just plug it into your PC printer port—plug your printer into the backpack. A simple solution from the backpack family of tape, diskette and hard drives.

backpack

MicroSolutions

Call Toll Free - 800.295.1214

132 W. Lincoln Hwy. DeKalb, Illinois 60115 • Telephone 815.756.3411 • FAX 815.756.2928

Our Family Tree is Branching Out!

274 BYTE APRIL 1995

THE BUYER'S MART

BAR CODE

Bar Code Printing Software
LabelWorks for Windows

➤ Prints all Popular Bar Code Types (19 Types)
➤ Desktop Publishing Features: WYSIWYG, Scalable Fonts, Rulers, Guides, Lines, Shapes, Page Zooms (25%-400%), Templates
➤ Rotates Text, Bar Codes, and Graphics
➤ Supports Windows Compatible Fonts
➤ Choose From Over One Hundred Popular Label Formats or Design Your Own
➤ Rich Text Support: Mix Styles, Types, & Sizes
➤ Automatically Prints Serial Numbers
➤ Imports & Exports Graphic Files: TIFF, GIFF, BMP, PCX, WPG, WMF, TARGA
➤ Supports Virtually all Windows Compatible Printers (PostScript, Laser, & Dot Matrix)
➤ 30-day Money-Back Guarantee, $295

CALL FOR FREE DEMO SOFTWARE

AMERICAN MICROSYSTEMS
2190 Regal Parkway, Euless, TX 76040
(800) 648-4452 (817) 571-9015 FAX (817) 685-6232

Inquiry 655.

BARCODE & MAG. STRIPE SYSTEMS

- **Keyboard Wedge** with HP Stainless Steel Wand/Mag. Stripe Reader **$249**
- **Keyboard Wedge** with SYMBOL LS2000 or SP400 Laser & Mag. Stripe Reader **$849**
- **Keyboard Wedge** with PSC QuickScan Laser/Mag. Stripe Reader **$699**
- **Software Wedge Decoder** with HP Stainless Steel Wand or Laser Scanner (DOS & WIN- RS-232 or parallel) **$189 +**
- All **Wedge Packages** include a Wand or Laser Holder
- **Mag. Stripe Encoder/Reader** (3 Trks) w/Software **$1099**
- **Printing Software** (DOS, WIN, UNIX...) **$149 +**
- **Portable Data Terminals** (128K– 4.2MB) **$599 +**
- **Complete POS System:** 486 40Mhz, **$1999** 4MB RAM, monitor, POS Software, SP212 Receipt Printer, M–S Cash Drawer, pole display, HP stainless steel wand and magnetic stripe reader with decoder
- **Application Software:** Inven, Asset, Tools, Time & Attend...
- **Radio Frequency** Terminals (spread spectrum/narrow band)
- Bar Code Printing Software (DOS) included with each purchase
- Made in the USA • 30 Day $$ Back • Spanish Dept. Avail. • Direct from Mfg.

BARCODE INTERNATIONAL SYSTEMS (BIS)
12140 Severn Way, Riverside, CA 92503 (909) 270-0016 Int'l
(800) 653-4252 US • (800) 219-5178 CAN • FAX (909) 270-0920

Inquiry 655.

CAD

Circuit Design Software for Windows
Easy-to-use schematic entry, PCB design, and simulation software, starting at $149 each. Complete PCB package with schematics, autorouter, and layout for 2-layer circuit boards, $399. Enhanced version with autoplacement, more symbol libraries, and up to 16 layers, $649. CAM file outputs.
Mental Automation, Inc.
5415 136th Place, SE-Bellevue WA 98006
(206) 641-2141 FAX (206) 649-0767 BBS (206) 641-2846

Inquiry 656.

CAD/CAM

CONTOURING MOTION CONTROL
FROM A PRINTER PORT!
Indexer LPT™ software **$249**
NEW VERSION 3 VISA/MC
• Controls up to six step motors simultaneously.
• Linear and Circular Interpolation.
• New features to accommodate machine control.
• Easy-to-use DOS device driver. Super Manual.
• CAD-CAM interface available.
Ability Systems Corporation, 1422 Arnold Ave.
Roslyn, PA 19001 (215) 657-4338
FAX: (215) 657-7815

Inquiry 657.

CD-ROM

CD-ROM TITLES

101 Best Games #1 or #2	$13	Family Dr. 3rd Ed.	$17
11 Mil Business PhBk	$18	Jump Raven	$32
70 Mil Home PhBk	$23	Mathematics Lib.	$13
Apogee's Greatest Games	$18	Middle School Suite	$18
Baseball Cards	$20	Pinball Dreams Deluxe	$20
Business Made Easy	$15	SI '94 Sports Almanac	$20

CALL/WRITE/FAX FOR FREE PRICE LIST
Over 500 titles

ALLTEX COMPUTERS PHONE: (800) 891-5134
3900 N. 1st, Suite 7 (915) 675-5134
Abilene, TX 79603 FAX: (915) 676-4360
MASTER CARD - VISA - COD - PREPAID

Inquiry 658.

We Buy, Sell & Trade
CD-ROMS & MEMORY CHIPS
Resellers Wanted
Call or write for a free product update
Consolidated CDROM, Inc.
515 67th Ave Philadelphia PA 19126 USA
+1-215-276-3657 / +1-215-276-3854 fax
1-800-8-CDROMS

Inquiry 659.

INTERNET on CDROM!

LINUX Developers Resource 3 CD Set $25
Slackware & more! Complete materials to install & run!

CICA Windows 3 CD Set $25
Thousands of plug & play windows programs! 250+ games!

INTERNET Tools CD-ROM $30
Networking tools & utilities for DOS & UNIX

BSDisc (NetBSD & FreeBSD) $35
Ready to use formats with install scripts

USENET 2 CD Set $25
comp.sources & alt. sources + many FAQ's from other groups

SOURCE CODE CD-ROM $30
4.4 BSD-Lite, XIIR6, MACH, Andrew Windowing

PERL & TCL/TK CD-ROM $35
Utility lang + command lang & toolkit for X-Windows

STANDARDS 2 CD Set $30
RFC's, IEN's, CCITT/ITU Bluebook, Windows Sockets

MS-DOS 2 CD Set (SimTel Archive) $25
10,000+ DOS programs, utilities, BBS software & docs

MC, VISA & AMEX   1-800-800-6613

tel: +1-602-526-9565
InfoMagic fax: +1-602-526-9573
P.O. Box 30370, Flagstaff, AZ 86003-0370 info@infomagic.com

Inquiry 660.

CD ROM TOWERS & JUKEBOX SERVERS FOR ALL OPERATING SYSTEMS!

No Device Drivers/ MSCDEX needed,
Complete Kit Networks CD Roms,
unlimited user license, DISCPORT.

"JES, NONE BETTER AT ANY PRICE"
Call NOW: 1 (800) 482-1866 305-271-0076

Inquiry 661.

EVERY CD ROM YOU'LL EVER NEED
TO CREATE EXCITING MULTIMEDIA
MULTIMEDIA MEGA BUNDLE
5 DISCS AT AN AMAZING PRICE!!!
GRAPHICS, AUDIO, MULTIMEDIA, MEGA I & II
~~$49.95~~ **$39.95** If you mention BYTE!

KNOWLEDGE MEDIA INC.
435 NUNNELEY RD., PARADISE, CA 95969
(800) 78 CD ROM, (916) 872-7487, (916) 872-3826 FAX

Inquiry 662.

CD-ROM

New and Updated CDROM Titles

Cica MS Windows CDROM, Thsnds of Windows prgrms..$29.95
Giga Games CDROM, Games for DOS/Windows...............$39.95
Space and Astronomy, Thsnds NASA images/data...........$39.95
C User Group Library, C source code Dec 93. ....$49.95
Simtel MSDOS CDROM, DOS Shareware/Freeware $29.95
QRZ Ham Radio CDROM, FCC Callsign Db & Shrwar......$29.95
Hobbes OS/2 CDROM, OS/2 Shareware/Freeware $29.95
Source Code CDROM, 650 Mb source, DOS/Unix............$39.95
Gutenberg Project, Literature and docs...........................$39.95
Linux Operating Sys, 386/486 OS, X11, full src.................$49.95
FreeBSD Operating Sys, Ver 1.0, krnl src, X/GNU $39.95
Libris Britannia, MSDOS Tech/Sci/Engineer.....................$69.95
X11R5/Gnu CDROM, Full src, SPARC binaries $39.95
Nebula for NeXTSTEP, Prgms for Intel NeXTSTEP $59.95
Ada Programming CDROM, Compilers, source, docs.........$39.95
Aminet CDROM, Amiga Shareware/Freeware $29.95
CDROM Caddies, Lifetime Guarantee $4.95

Top quality CDROMs. 100% satisfied or full refund.

WALNUT CREEK CDROM
4041 Pike Lane, Ste D-212, Concord, CA 94520
1-800-786-9907 Visa/MC AMEx, Fax: 1-510-674-0821

Inquiry 663.

COMMUNICATIONS

Frame Relay, X.25, BSC, HDLC, SDLC
Use our rock solid, compliant, inexpensive and robust synchronous products for your PC project. On board protocol support reduces PC overhead.
• Support for MS-DOS, Windows, Unix, OS/2, Netware and others.
• CCITT and ISO compliant X.25, HDLC.
• Frame Relay blanket certified for any application.
• Test and datascope programs for easy debugging.
Sangoma Technologies Inc.
Tel 1-800-388-2475 905-474-1990
Fax 905-474-9223

COMPUTER BOOKS

COMPUTER BOOKS ONLINE
Easy online ordering at our Internet, CompuServe, & NIFTYServe stores. 15% off most books from 140 publishers for computer professionals. GO CBK, www.compubooks.com/books.html, or ftp info.txt from ftp.compubooks.com. Worldwide shipping. E-mail to info@compubooks.com for information.
CompuBooks
RR1 Box 271D 512-321-9652
Cedar Creek TX 78612 USA Fax 512-321-4525
Order line 800-880-6818

COMPUTER INSURANCE

INSURES YOUR COMPUTER
SAFEWARE Computerowner's coverage provides replacement of hardware, media and purchased software. As little as $49 a year covers accidents, theft, power surges and more. One call does it all.
1-800-800-1492
SAFEWARE, The Insurance Agency Inc.
PO Box 02211, 2929 N. High St., Columbus, OH 43202
Now available in Ontario!!!

Inquiry 664.

COMPUTER TELEPHONY

Write your own Voice Mail software
Create advanced, multi-line voice/fax mail systems with VOS™, the Editors' Choice. Complete language for computer/telephone integration. VOS supports databases, LANs, voice recognition, text-to-speech, conferencing, switching and more.
Call for free booklet *Get into Interactive Voice.*
Parity Software
US: 415-989-0330 fax: 415-989-0441
Europe: +45-3940.8803 fax: +45-3940-7803

THE BUYER'S MART

LANS

Little Big LAN
The most flexible network
- Peer to Peer LAN to 250 nodes
- $75 total software cost, not per node!
- Link via serial, parallel, or Modems
- Also via Ethernet or Arcnet, or mix!
- Typically only 40k of RAM

Information Modes
817-387-3339 / P.O. Drawer F, Denton TX 76202
Fax 817-382-7407   Orders 800-628-7992

Inquiry 680.

LASER CHECKS

LASER CHECKS
PERSONALIZED WITH YOUR NAME AND LOGO
IN SAFETY BACKGROUND
CALL (714) 773-5811 OR TOLL FREE
1-800-252-6427   Fax 1-800-439-0158
FOR SAMPLE (ETC.)
:MAGNETIC ENCODING:
MICR-COMP, Inc.
689 S. State College Blvd., Suite A
Fullerton, CA 92631

Inquiry 681.

PROGRAMMERS TOOLS

DON'T DEVELOP NETWORK APPLICATIONS
without **NPPC**, a powerful library subroutine
package with the simplest API available anywhere!
Window and DOS, IPX or NetBIOS. Use with any
language supporting C subroutine calls: C, C++,
Visual Basic, FoxPro, etc. 30 day $ back
guarantee. Download a free demo from our BBS.
NPPC for Windows or DOS – $195 w/source $395.
SYNERGY SOLUTIONS, INC.
2150 South Country Club , Suite 1, Mesa, AZ 85210
(602) 545-9797 FAX (602) 545-9827 BBS (602) 545-0232

Inquiry 682.

PUBLICATIONS

SCSI MYSTERIES UNVEILED
Whether a neophyte or an expert, find what you
need in the "SCSI Bench Reference", "SCSI
Tutor", and "SCSI Encyclopedia".

ENDL Publications
408-867-6642   Fax: 408-867-2115
Receive immediate information via FaxAccess.
Dial 408-741-1600 from the handset,
order #5 and follow instructions on voice menu.

Inquiry 683.

SECURITY

EVERLOCK – EVERKEY
Software and hardware based copy protection.
Features include – encryption, serialization,
compression, remotely resettable date and
execution limits, user data flags and more!
Call today and ask about our low cost Trial Kits
and free demo disk.
Az-Tech Software, Inc.
201 East Franklin St., Richmond, MO 64085
(800) 227-0644   (816) 776-2700   FAX (816) 776-8398

Inquiry 684.

THE ULTIMATE SOFTWARE SECURITY
- STOPCOPY family – UNCOPIABLE copy protection
- STOPVIEW software encryption
- NETLIMIT network license metering
- DOS, Windows, Macintosh, OS/2, support
- No source code changes required - for ANY of our products in ANY environment
- Our products destroy ALL of our competition
- Call for FREE demo disk, or to discuss our products' MANY options

BBI Computer Systems, Inc.
14105 Heritage Lane, Silver Spring, MD 20906
800/TRY-ABBI • 800/879-2224 • 301/871-1094 • FAX:301/460-7545

Inquiry 685.

SECURITY

CRYPKEY SOFTWARE LICENSING SYSTEM
"Hardware key protection without the hardware key"
CrypKey is a **software protection tool**, offering
- complete security from any disk copy program
- complete compatibility wtih any MS-DOS or MS WINDOWS based machine
- complete invisibility – no disk key, no hardware key, less support calls
- instant disaster recovery
- protect programs distributed on CD-ROM, BBS, or Internet!

CrypKey is a **sales tool,** allowing you to sell your program
- by increments – enable the options the customer purchased
- by number of runs – e.g., sell 100 calculations for $100.00
- by time period – e.g., lease or demo your program for 60 days

CrypKey uses a numeric key that can be transmitted by phone,
fax, or email. Sell your customers more options, more copies,
more time or more runs instantly, just by making a telephone
call. (great for overseas customers or distributors) CrypKey is
produced by Kenonic Controls Ltd. - engineering and software
since 1972.

Kenonic Controls Limited
7175–12th Street South East
Calgary, Alberta, Canada T2H 2S6
(403) 258-6200 • fax: (403) 258-6201

Inquiry 686.

Cop's CopyLock II
Professional software protection with
TRUE Machine Install. Option Board safe.
DOS, OS2, Networks, Windows, Trace 3000.
DialCOPS Access Control for mass
distribution via CD-ROM or Internet.
Known and used world-wide since 1984.
LINK Data Security
Int'l: + 45 3123-2350 Fax: + 45 3123-8448

KEY-LOK™ SECURITY
Piracy survival 11 years proves effectiveness
of powerful multilayered security. Algorithmic
response. Programmable memory. Econom-
ical. Transparent to PARALLEL/SERIAL port,
Counters/Real-Time-Clock. Multi-product/
feature licensing. DOS/NT/UNIX/OS2.
Access control systems and disk drive locks.
MICROCOMPUTER APPLICATIONS
3167 E. Otero Circle, Littleton, CO 80122
1-800-4KEYLOK (303) 770-1917 FAX: (303) 770-1863

Inquiry 687.

SERIAL ANALYZER

COM-Watch PROFESSIONAL
Use your PC (even your notebook) to trouble-
shoot RS-232, RS-422 or RS-485 communications
at a fraction of the cost of purchasing specialized
equipment. This menu driven software lets you
monitor bi-directional communications at speeds
to 115,200 baud. Package includes RS-232 cable
adapter.
CER • P.O. Box 48 • Amherst, NH 03031
Tel: (603) 673-8110 • Fax (603) 673-7892

Inquiry 688.

SOFTWARE

CAD to NC CODE Instantly!
with FastPATH™
Path complex parts to whole nests!
Hands free, Automatic & Intelligent
Dos based, Suits all NC profilers.
**Call FastCAM at (303) 667-5059
or Fax: (303) 667-1990**
Ft. Collins, Colorado, U.S.A.

Inquiry 689.

SOFTWARE PACKAGING

FREE SOFTWARE PACKAGING CATALOG
Everything you will need to Package, Distribute, and Ship Your
Software!! From manuals and binders to mailers and shippers

LABELS • LABELS • LABELS
For your diskettes, plain or custom printed
dot matrix or laser printer . . . free samples
•••FREE CATALOG•••
Hice & Associates
8586 Monticello Dr., West Chester. OH 45069
Phone/Fax: 513-779-7977

Inquiry 690.

SOFTWARE/BUSINESS

DATA ENTRY SOFTWARE
Full featured, heads-down data entry
with two-pass verification, edit language,
operator stats, much more! Designed for
the PS/2®, PC, XT, AT or compatibles.
PCs from $395   LAN version available
FREE 30 day trial

Computer Keyes	Tel: 206/776/6443
21929 Makah Rd.,	Fax: 206/776-7210
Woodway, WA 98020	USA: 800/356-0203

Inquiry 691... (duplicate?)

SOFTWARE/ENGINEERING

Analog/Digital Simulation!!
- DOS, Alpha, Macintosh
- Windows & WIndows NT
- IsSPICE4 Real Time SPICE
- Analog/Mixed SImulation
- Schematic Entry
- New AHDL Modeling Kit!!
- Model Libraries, RF, Power
- More Than 5000 parts
- Waveform Analysis
- Full SPICE programs starting at $95. Complete systems, $595-$2595

P.O. Box 710 San Pedro, CA 90733-0710
(310)833-0710, FAX (310)833-9658   **intusoft**
Call for your Free Demo and information kit.

Inquiry 691.

SAUNA: 3D THERMAL ANALYSIS
● Models: PCBs, stacked plates, heatsinks, multiboard
enclosures. ● All heat transfer modes: convection,
radiation, conduction ● Interactive menu-driven
● Thermal parameters library ● Fast "What if":
dimension, mat'l, finish, analyses ● Easy to learn & use
● IBM PC & Macintosh II
Call or FAX for free evaluation program
Tatum Labs, Inc.
1287 N. Silo Ridge Drive, Ann Arbor, MI 48108
313-663-8810   FAX 313-663-3640

Inquiry 692.

SOFTWARE/GRAPHICS

New Version!
AccuSoft Image Format Library 5.0

Programmers: Add support for 36 raster
file formats instantly!

TIFF, JPEG, PCX, TARGA, DIB, DCX, GIF, BMP,
WMF, PICT, WPG, EPS, Group 3, Group 4
New Formats: Photo CD, PhotoShop, ASCII,
KoFax, RLE, LaserData, CALS, ATT, CLP, XWD,
IMG, IFF, SUN, XBM, ICO, IOCA, CX2, XPM,
CUT, Brooktrout, MAC, MSP.

*Guaranteed to read all raster images in existence
in the listed formats!*

✳ Import, export, scanning, conversion, compression
✳ Printing, display, image processing
✳ Supports all languages
✳ Fax formats and multi-page images
✳ Rotate, zoom, scale, color reduction
✳ Thumbnails, sharpen, special-effects
✳ Windows, NT, WinPro Gold 32, VBX, VBX32
✳ Watcom, OS/2, MAC, UNIX, Clipper, FoxPro

AccuSoft Corp. Call 800-525-3577
Two Westborough Business Park   Westborough, MA 01581 USA
TEL (508) 898-2770   FAX (508) 898-9662

Inquiry 693.

To order products or request FREE information, call advertisers directly or send in the response card by mail or fax! **Let them know you saw it in BYTE!**

Inquiry No.		Page No.	Phone No.
	A		
286	ADAPTEC	264	408-945-8600
229-230	ADONTEC GMBH	273	+49-7043-40440**
237	ADVANCED ENGINEERING CONCEPTS	270	310-379-1189
212-213	ADVANCED IMAGE COMMNCTNS	269	800-209-2418
*	ADVANCED MICRO DEVICES (N.A.)	12-13	800-222-9323
235-236	AEROCOMM WIRELESS INC	266	800-492-2320
128-129	AERONICS INC	39	512-258-8040
61	ALADDIN KNOWLEDGE SYSTEMS	87	800-223-4277
563	ALTEX ELECTRONICS	260NE 2-3	800-531-5369
564	ALTEX ELECTRONICS	260SO 2-3	800-531-5369
	AMERICA ONLINE INC (N.A.)	180A-B	800-827-6364 PC#11069
174	AMERICAN ADVANTECH	268	800-800-6889
*	AMERICAN LUNG ASSOC	260NE 1	
*	AMERICAN LUNG ASSOC	260SO 1	
*	AMERICAN POWER CONV	16A-B	
62	AMERICAN POWER CONV	16-17	800-800-4APC dept A2
291-292	ANDATACO	162F-G	800-334-9191
*	ANTHRO CORP	80	800-325-3841
561-562	APE PTACEK ENGINEERING GMBH	32IS 9	+49-89-604395**
534-535	ARISTO GRAPHIC SYSTEMS	32IS 17	+49-40-54747-111
409	ARTECON	88DM 9	800-872-2783
573-574	ARTISOFT UK (INT'L)	157	+44(0)753 55 49 99
158	AVALAN TECHNOLOGY	261	508-429-3179**
575	AXIS COMMUNICATIONS (INT'L)	141	+46 46 140500
	B		
299	BADGER COMPUTER	162H	800-3-BADGER
*	BBN SOFTWARE PRODUCTS (N.A.)	80A-B	
137	BBN SOFTWARE PRODUCTS (N.A.)	81	800-331-2266 ext 130
*	BEAME & WHITESIDE S/W	32IS 28	919-831-8989
529	BEAME & WHITESIDE S/W	32IS 29	919-831-8989
260	BEAME & WHITESIDE S/W (N.A.)	157	800-216-8450
63	BELL & HOWELL INC (INT'L)	79	708-675-7600
63	BELL & HOWELL INC (N.A.)	79	800-SCAN-494
450	BIX	285	800-695-4775
501-502	BOCA RESEARCH INC	32IS 30	407-997-6227
223-224	BOXLIGHT CORP	270	800-762-5757
140	BTG	198	703-876-1920**
144	BUFFALO INC	261	800-345-2356
*	BYTE ANNIVERSARY POLL (U.S.)	170	
*	BYTE BACK ISSUES	260NE 4	
*	BYTE BACK ISSUES	260SO 4	
*	BYTE BACK ISSUES (INT'L)	225	
*	BYTE EDITORIAL SURVEY	82	
*	BYTE EURODECK (CD)	170	
*	BYTE MOVING	32IS 20	
*	BYTE NETWORKING	210	
65	BYTE ON CD ROM	184	
*	BYTE READER (INT'L)	174	
*	BYTE REPRINTS (INT'L)	219	603-924-2525
*	BYTE SUB MESSAGE	138	
*	BYTE SUB MESSAGE	32IS 10	
	C		
162	CADSOFT COMPUTER INC	257	407-274-8355
130	CALIFORNIA PC PRODUCTS INC	229	800-394-4122
559-560	CALLUNA TECHNOLOGY LTD	32IS 22	+44(0)1592 630920**
281	CASEWISE LTD	187	800-893-8398
503	CHERRY MIKROSCHALTER GMBH	32IS 12-13	+49-9643-18-206
506	COBALT BLUE	32IS 24	404-518-1116
557-558	COMBYTE INC	32IS 25	303-229-0660
239	COMPAQ (N.A.)	127	800-392-8883
504-505	COMPEX INC (INT'L)	127	714-630-7302
548	COMPEXPO / COMPFAIR	32IS 26	+36-1-117-0436**
283	COMPUTER ASSOC-REALIA	57	800-434-REAL dept 26500
289	COMPUTER ASSOC-UNICENTER	115	800-225-5224 dept 10500
145	COMPUTER DISCT WAREHSE	240-241	800-959-4CDW
*	COMPUTER MUSEUM	178	
*	COMPUTER PROFESSIONALS' BOOK SOC (N.A.)	196A-B	
*	COMPUTER PROFESSIONALS' BOOK SOC (N.A.)	197	717-794-2191
507	COMPUTER QUICK	32IS 31	415-861-8330
146	COMPUTERLANE UNLTD	248	800-526-3482
284	CONNER PERIPHERALS (N.A.)	221	800-6-CONNER
*	CONSENSEYS CORP	88DM 4	800-388-1896
175	CONTROL VISION	270	316-231-6647
*	COPIA INTERNATIONAL LTD	38	708-682-8898

Inquiry No.		Page No.	Phone No.
66	COREL CD CREATOR	27	613-728-3733 ext 28
222	CR SYSTEMS	271	800-732-3664
115	CREATIVE LABS INC	59	800-998-5227
240-241	CTX INTERNATIONAL INC	149	909-595-6293**
67-68	CURTIS INC	182	612-631-9512
147-148	CYBEX CORP	243	205-430-4030**
149-150	CYBEX CORP	255	205-430-4030**
508-509	CYBEX CORP (INT'L)	CIII	205-430-4030**
	D		
531-532	DAEWOO (INT'L)	11	+31-2503-37214**
163	DALLAS SEMICONDUCTOR	262	800-258-5061
233-234	DASYTEC	272	800-731-5015
264	DATA FOCUS (INT'L)	145	703-818-1532**
264	DATA FOCUS (N.A.)	145	800-637-8034
122-123	DATA STORAGE MARKETING INC (N.A.)	36	800-543-6090
267-268	DATADISC	155	800-DATADISC
151	DATALUX CORP	242	800-DATALUX
*	DATAPRO (INT'L)	116-117	800-328-2776
271-272	DATAPRODUCTS	162D	818-887-8000
69	DATASTORM TECH (N.A.)	219	800-315-3282
565	DECUS / TALLEY MANAGEMENT GROUP	260SO 1	
*	DELL COMPUTER CORP (N.A.)	CIII	800-934-5259
*	DELL COMPUTER CORP (N.A.)	CIV	800-876-1190
*	DELL COMPUTER CORP (N.A.)	CIV	800-873-5533
70	DELPHI INTERNET SERVICES	231	800-695-4005
243	DELRINA FORMFLOW	135	800-268-6082
242	DELRINA WINFAX PRO	133	800-268-6082
71-72	DELTEC / NSSI	96	800-335-8321
414-415	DIGIBOARD	88DM 16	800-551-4794
287	DIGITAL EQUIPMENT CORP	106-107	800-STORWORK
*	DIGITAL INFO SERVICE (EUROPE)	32IS 21	+41-22-7910885**
555-556	DIGITAL INFO SERVICE (EUROPE)	32IS 23	+41-22-7910885**
549	DISTINCT CORP	32IS 6	408-366-8933
152	DISTRIBUTED PROCESSING TECH	250	800-322-4378
577	DR HUGGLE & PARTNER GMBH (INT'L)	197	+49-241-403117**
	E		
202	EDUCALC	270	800-713-6525
214	EMATEK GMBH	272	+49 221 529666
572	EMATEK GMBH (INT'L)	97	+49 221 529666
544-545	ERGOTRON EUROPE	32IS 7	+31 20 696.60.65
553-554	EUROSOFT TECHNOLOGY	32IS 32	+44 1635 582 660
512-513	EUTRON	32IS 8	+39 35 201003
73-74	EXABYTE CORP	19	800-EXABYTE
	F		
542-543	F & H SIMULATIONS (INT'L)	69	+31 13 427 516**
514-515	FAST HARDLOCK	32IS 1	+49-89-539800-20
538-539	FAST MULTIMEDIA	32IS 5	+49-89-50206-199**
141	FIELDWORKS	228	612-947-0856
527	FIRST INTERNATIONAL COMP	32IS 2	+886-2-718-2782**
153-154	FIRST SOURCE INT'L	252	714-448-7750
576	FISKARS POWERS SYS (INT'L)	98	+358 0 452661
75	FRAME TECHNOLOGY (N.A.)	97	800-U4FRAME ext 618
217	FUZIWARE, INC	273	800-472-6183
	G		
176	GAGE APPLIED SCIENCES INC	268	514-337-6893
*	GATEWAY 2000	49	800-846-2058
*	GATEWAY 2000	51	800-846-2058
77	GLENCO ENGINEERING (INT'L)	103	708-808-0300
76	GLENCO ENGINEERING (N.A.)	103	800-562-2543
200-201	GRANITE DIGITAL	269	510-471-6442
516	GREY MATTER LTD (INT'L)	36	+44-(0)364-53071**
	H		
177	HOOLEON CORP	269	520-634-7515
550-551	HYPERSYSTEMS	32IS 20	+39-11-434-2350
	I		
269	IBM (N.A.)	112-113	800-IBM-6676 ext 621
412-413	ICL (EMBLA)	88DM 8	703-648-3326**
133-134	ICONOVEX INC (N.A.)	85	800-943-0292
296	IEM INC	88DM 23	800-321-4671
207	IERC	265	818-842-7277
78	INTEGRAND RESEARCH	226	209-651-1203
416	INTEGRIX INC	88DM 11	805-375-1055
566	INTEGRIX INC (INT'L)	161	805-375-1055
*	INTEL CORP (N.A.)	32A-D	800-538-3373
244-245	INTERGRAPH	129	800-763-0242

Inquiry No.		Page No.	Phone No.
400-401	INTERGRAPH / SOFTWARE DIVISION	88DM 19	800-345-4856
552	INTERSIS AUTOMACAO	32IS 18	+351 441 2247
178	IO TECH	269	216-439-4091
113-114	ITERATED SYSTEMS	88	800-437-2285
	J		
155	JAMECO ELECTRONICS	253	800-831-4242
*	JDR MICRODEVICES	263	800-538-5000
215-216	JNL COMPUTER PRODUCTS	270	510-353-0262
	K		
179	KILA	268	303-444-7737
246-247	KINGSTON TECHNOLOGY	147	714-435-2667
402-403	KINGSTON TECHNOLOGY	88DM 21	714-435-1850
518	KUO FENG CORP (INT'L)	121	+886-2-754-8498**
	L		
166-167	L A TRADE	246	800-433-3726
261	LANSOURCE TECHNOLOGY	162	416-535-3555
290	LEXMARK	119	800-891-0399 ext 230
301-302	LOCUS COMPUTING CORP (INT'L)	170	+44-1442-871774
519	LOGIC PROGRAMMING ASSOC	32IS 14	800-949-7567
125	LOGICIELS ET SERVICES DUHEM	228	+33149700455
	M		
520	MAGIC / MSE (INT'L)	71	+972-3-751-1901**
225	MAILER'S SOFTWARE	271	800-800-MAIL
533	MANNESMANN TALLY	32IS 15	+44 (0) 1734 788711
*	MCGRAW HILL NRI (N.A.)	212A-B	
135-136	MEDIA ON (N.A.)	174	800-552-7835
220-221	MENAI CORP	273	800-GAMELON
161	MICRO 2000	251	800-864-8008
156-157	MICRO SOLUTIONS COMP PROD	247	800-295-1214
194-195	MICROCAL SOFTWARE INC	273	800-969-7720
536-537	MICRODATA SYSTEM SRI	32IS 18	+39 1 87 988 460
277	MICRO-INTERNATIONAL INC	162E	800-967-5667
116	MICRON COMPUTER	CII,1	208-465-3434
248-249	MICROPOLIS CORP (N.A.)	116-117	800-395-3748
*	MICROSOFT CORP	15	
180	MICROSTAR LABORATORIES	269	206-453-2345
81	MICROWAY	222	508-746-7341
526	MINICOM / CLASSNET VIDEO	32IS 10	+972-2-518593
253-254	MINUTEMAN	120	800-238-7272
521-522	MINUTEMAN	32IS 16	214-446-7363
280	MIPS DATALINE AMERICA INC	162B	800-898-8560
300	MITSUBISHI MOTOR SALES OF AMERICA (N.A.)	138A-D	800-55MITSU
250	MKS / MORTICE KERN SYSTEMS	118	519-884-2251
	N		
82-83	NANAO USA CORP (N.A.)	71	310-325-5202
181	NATIONAL INSTRUMENTS	272	512-794-0100
*	NEKOTECH	88DM 7	714-580-0055
285	NEVADA COMPUTER	258	800-982-2925
139	NEXGEN	189	408-435-0202
84-85	NORTON-LAMBERT	63	800-964-6767
86	NSTL	235	610-941-9600
87	NUMBER NINE COMPUTER SYS	31	800-GET-NINE
	O		
131	OBJECTS INC	193	508-777-2800
517	OLIVETTI S.P.A. (INT'L)	12-13	+39-2-45361
273-274	OMNICOMP GRAPHICS CORP	162C	713-464-2990
528	ON TIME MARKETING	32IS 24	+49-40-437472
251-252	OPTIQUEST	162A	909-468-3750
*	ORACLE CORP (N.A.)	111	800-633-0542 ext 4990
170-171	ORCHESTRA MULTISYSTEMS	264	800-237-9988
*	OSBORNE MCGRAW-HILL	40-41	800-822-8158
*	OSBORNE MCGRAW-HILL	260	800-822-8158
*	OSBORNE MCGRAW-HILL (INT'L)	112-113	800-822-8158
	P		
168	PACIFIC COAST MICRO	254	619-581-6040
88-89	PC POWER & COOLING	45	800-722-6555
169	PC'S COMPLEAT	244-245	508-624-6400
90	PERSOFT INC	223	800-368-5283
182	PERSONAL TEX	272	800-808-7906
203-204	PIKA TECHNOLOGIES	266	613-591-1555
91-92	PINNACLE MICRO	7	714-727-3300
93	PKWARE INC	104	414-354-8699
94	PKWARE INC	180	414-354-8699
126-127	PLEXTOR	21	800-886-2683

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today.
For quickest response, fax to 1-416-620-9790!

Product Category Information

To receive information for an entire category, circle the appropriate number on the adjacent card.

Hardware

Accessories/Supplies	1
Add-in Boards	2
Bar Coding	3
Communications/Networking	4
Computer Systems	5
Data Acquisition	6
Diagnostic Equipment	53
Disks & Optical Drives	7
Diskettes/Duplicators	8
Fax Boards/Machines	9
Graphics Tablets/Mice/ Pen Input	10
Keyboards	11
LAN Hardware	12
Laptops & Notebooks	13
Mail Order	14
Memory/Chips/Upgrades	15
Miscellaneous Hardware	16
Modems/Multiplexors	17
Monitors & Terminals	18
Multimedia/CD-ROM	19
PCMCIA	57
Printers/Plotters	20
Programmable Hardware	21
RAID Drive Arrays	56
Scanners/OCR/Digitizers	22
Security	52
Tape Drives	23
UPS/Power Management	24
Voice Technology	55

Software

Business	25
CAD/CAM	26
Communications/Networking	27
Data Acquisition	28
Database	29
Educational	30
Engineering/Scientific	31
Entertainment	32
Graphics	33
Macintosh	34
Mail Order	35
Mathematical/Statistical	36
Miscellaneous Software	37
On-Line Services	38
Operating Systems	39
OS/2	54
Programming Languages/ Tools	40
Security	41
Shareware	42
Software Duplication	43
Spreadsheets	44
Unix	45
Utilities	46
Windows	47
Word Processing	48

General

Books/Publications	49
Recruitment	50
Miscellaneous	51

Fill out this coupon carefully. *Please Print.*

Name _____ Title _____

Company _____ Address _____

City _____ Country _____ Code _____

Phone _____ Fax _____

A. My job function is best classified as (check one):
1 [ ] Company Management
2 [ ] IS/MIS/IT Management
3 [ ] Systems Engineering/Integration
4 [ ] Systems/Networking Consulting
5 [ ] Departmental Management (non-IS/MIS)
6 [ ] Technical Services Support
7 [ ] Other (please describe)

B. What is your organization's primary business activity at this location (check one):
8 [ ] Business Services (Finance, Banking, Insurance, Healthcare, Professional)
9 [ ] Commerce/Industry (Retail, Wholesale, Construction, Mining, Manufacturing, Transportation)
10 [ ] Reseller/OEM (VAR, VAD, Systems/Network Integrator, Computer Product Manufacturer)
11 [ ] Government (Federal, State, Municipal, Military)
12 [ ] Computer Services (Support, Training, Consulting)

13 [ ] Other (please specify):
_____

C. Please indicate which specific fields of interest are important to you so that BYTE may serve you with supplemental information that best meets your needs (select all that apply):
14 [ ] UNIX and Workstations
15 [ ] Networking (LANs, WANs, and Telecommunications)
16 [ ] Multimedia
17 [ ] Reselling/Systems Integration
18 [ ] Applications Development
19 [ ] Other (explain):
_____

D. My responsibilities require that I be involved with the following operating system environments (check all that apply):
20 [ ] DOS
21 [ ] OS/2
22 [ ] Mac/OS

23 [ ] UNIX (any, including Solaris)
24 [ ] Windows
25 [ ] Windows/NT
26 [ ] NetWare
27 [ ] Other (please describe):
_____

E. The number of employees at my location and company-wide area (check one in each column):

	At My Location	Company-Wide
10,000 or more	28 [ ]	33 [ ]
5,000 to 9,999	29 [ ]	34 [ ]
1,000 to 4,999	30 [ ]	35 [ ]
100 to 999	31 [ ]	36 [ ]
Under 100	32 [ ]	37 [ ]

April 1995   90 97 99
Valid until June 30, 1995

Inquiry Numbers 1–187

```
 1   2   3   4   5   6   7   8   9  10  11  12  13  14  15  16  17
18  19  20  21  22  23  24  25  26  27  28  29  30  31  32  33  34
35  36  37  38  39  40  41  42  43  44  45  46  47  48  49  50  51
52  53  54  55  56  57  58  59  60  61  62  63  64  65  66  67  68
69  70  71  72  73  74  75  76  77  78  79  80  81  82  83  84  85
86  87  88  89  90  91  92  93  94  95  96  97  98  99 100 101 102
103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119
120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136
137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153
154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170
171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187
```

Inquiry Numbers 188–374

```
188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204
205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221
222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238
239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255
256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272
273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289
290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306
307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323
324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340
341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357
358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374
```

Inquiry Numbers 375–561

```
375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391
392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408
409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425
426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442
443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459
460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476
477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493
494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510
511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527
528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544
545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561
```

Inquiry Numbers 562–748

```
562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578
579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595
596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612
613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629
630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646
647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663
664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680
681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697
698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714
715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731
732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748
```

Inquiry Numbers 749–935

```
749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765
766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782
783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799
800 801 802 803 804 805 806 807 808 809 810 811 812 813 814 815 816
817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833
834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850
851 852 853 854 855 856 857 858 859 860 861 862 863 864 865 866 867
868 869 870 871 872 873 874 875 876 877 878 879 880 881 882 883 884
885 886 887 888 889 890 891 892 893 894 895 896 897 898 899 900 901
902 903 904 905 906 907 908 909 910 911 912 913 914 915 916 917 918
919 920 921 922 923 924 925 926 927 928 929 930 931 932 933 934 935
```

Inquiry Numbers 936–1122

```
 936  937  938  939  940  941  942  943  944  945  946  947  948  949  950  951  952
 953  954  955  956  957  958  959  960  961  962  963  964  965  966  967  968  969
 970  971  972  973  974  975  976  977  978  979  980  981  982  983  984  985  986
 987  988  989  990  991  992  993  994  995  996  997  998  999 1000 1001 1002 1003
1004 1005 1006 1007 1008 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1020
1021 1022 1023 1024 1025 1026 1027 1028 1029 1030 1031 1032 1033 1034 1035 1036 1037
1038 1039 1040 1041 1042 1043 1044 1045 1046 1047 1048 1049 1050 1051 1052 1053 1054
1055 1056 1057 1058 1059 1060 1061 1062 1063 1064 1065 1066 1067 1068 1069 1070 1071
1072 1073 1074 1075 1076 1077 1078 1079 1080 1081 1082 1083 1084 1085 1086 1087 1088
1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105
1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122
```

Inquiry Numbers 1123–1309

```
1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139
1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150 1151 1152 1153 1154 1155 1156
1157 1158 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1170 1171 1172 1173
1174 1175 1176 1177 1178 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1190
1191 1192 1193 1194 1195 1196 1197 1198 1199 1200 1201 1202 1203 1204 1205 1206 1207
1208 1209 1210 1211 1212 1213 1214 1215 1216 1217 1218 1219 1220 1221 1222 1223 1224
1225 1226 1227 1228 1229 1230 1231 1232 1233 1234 1235 1236 1237 1238 1239 1240 1241
1242 1243 1244 1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258
1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275
1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292
```

Inquiry Numbers 1310–1489

```
1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1320 1321 1322 1323 1324 1325 1326
1327 1328 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1340 1341 1342 1343
1344 1345 1346 1347 1348 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360
1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377
1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394
1395 1396 1397 1398 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411
1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428
1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445
1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462
1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479
1480 1481 1482 1483 1484 1485 1486 1487 1488 1489
```

BYTE

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today.
For quickest response, fax to 1-416-620-9790!

BYTE

TNT MAILFAST
FRG/BOS/700252/MO1
P.O. Box 66
Hounslow
TW5 9RT
England

Place
Postage
Here

BYTE

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today. For quickest response, fax to 1-416-620-9790!

Product Category Information

To receive information for an entire category, circle the appropriate number on the adjacent card.

Hardware

Accessories/Supplies	1
Add-in Boards	2
Bar Coding	3
Communications/Networking	4
Computer Systems	5
Data Acquisition	6
Diagnostic Equipment	53
Disks & Optical Drives	7
Diskettes/Duplicators	8
Fax Boards/Machines	9
Graphics Tablets/Mice/ Pen Input	10
Keyboards	11
LAN Hardware	12
Laptops & Notebooks	13
Mail Order	14
Memory/Chips/Upgrades	15
Miscellaneous Hardware	16
Modems/Multiplexors	17
Monitors & Terminals	18
Multimedia/CD-ROM	19
PCMCIA	57
Printers/Plotters	20
Programmable Hardware	21
RAID Drive Arrays	56
Scanners/OCR/Digitizers	22
Security	52
Tape Drives	23
UPS/Power Management	24
Voice Technology	55

Software

Business	25
CAD/CAM	26
Communications/Networking	27
Data Acquisition	28
Database	29
Educational	30
Engineering/Scientific	31
Entertainment	32
Graphics	33
Macintosh	34
Mail Order	35
Mathematical/Statistical	36
Miscellaneous Software	37
On-Line Services	38
Operating Systems	39
OS/2	54
Programming Languages/ Tools	40
Security	41
Shareware	42
Software Duplication	43
Spreadsheets	44
Unix	45
Utilities	46
Windows	47
Word Processing	48

General

Books/Publications	49
Recruitment	50
Miscellaneous	51

Fill out this coupon carefully. *Please Print.*

Name _____ Title _____

Company _____ Address _____

City _____ Country _____ Code _____

Phone _____ Fax _____

A. My job function is best classified as (check one):
1 [ ] Company Management
2 [ ] IS/MIS/IT Management
3 [ ] Systems Engineering/Integration
4 [ ] Systems/Networking Consulting
5 [ ] Departmental Management (non-IS/MIS)
6 [ ] Technical Services Support
7 [ ] Other (please describe)

B. What is your organization's primary business activity at this location (check one):
8 [ ] Business Services (Finance, Banking, Insurance, Healthcare, Professional)
9 [ ] Commerce/Industry (Retail, Wholesale, Construction, Mining, Manufacturing, Transportation)
10 [ ] Reseller/OEM (VAR, VAD, Systems/Network Integrator, Computer Product Manufacturer)
11 [ ] Government (Federal, State, Municipal, Military)
12 [ ] Computer Services (Support, Training, Consulting)

13 [ ] Other (please specify):
_____

C. Please indicate which specific fields of interest are important to you so that BYTE may serve you with supplemental information that best meets your needs (select all that apply):
14 [ ] UNIX and Workstations
15 [ ] Networking (LANs, WANs, and Telecommunications)
16 [ ] Multimedia
17 [ ] Reselling/Systems Integration
18 [ ] Applications Development
19 [ ] Other (explain):
_____

D. My responsibilities require that I be involved with the following operating system environments (check all that apply):
20 [ ] DOS
21 [ ] OS/2
22 [ ] Mac/OS

23 [ ] UNIX (any, including Solaris)
24 [ ] Windows
25 [ ] Windows/NT
26 [ ] NetWare
27 [ ] Other (please describe):

E. The number of employees at my location and company-wide area (check one in each column):

	At My Location	Company-Wide
10,000 or more	28 [ ]	33 [ ]
5,000 to 9,999	29 [ ]	34 [ ]
1,000 to 4,999	30 [ ]	35 [ ]
100 to 999	31 [ ]	36 [ ]
Under 100	32 [ ]	37 [ ]

April 1995   90 97 99
Valid until June 30, 1995

Inquiry Numbers 1–187

```
  1   2   3   4   5   6   7   8   9  10  11  12  13  14  15  16  17
 18  19  20  21  22  23  24  25  26  27  28  29  30  31  32  33  34
 35  36  37  38  39  40  41  42  43  44  45  46  47  48  49  50  51
 52  53  54  55  56  57  58  59  60  61  62  63  64  65  66  67  68
 69  70  71  72  73  74  75  76  77  78  79  80  81  82  83  84  85
 86  87  88  89  90  91  92  93  94  95  96  97  98  99 100 101 102
103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119
120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136
137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153
154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170
171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187
```

Inquiry Numbers 188–374

```
188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204
205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221
222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238
239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255
256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272
273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289
290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306
307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323
324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340
341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357
358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374
```

Inquiry Numbers 375–561

```
375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391
392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408
409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425
426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442
443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459
460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476
477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493
494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510
511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527
528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544
545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561
```

Inquiry Numbers 562–748

```
562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578
579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595
596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612
613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629
630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646
647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663
664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680
681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697
698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714
715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731
732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748
```

Inquiry Numbers 749–935

```
749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765
766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782
783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799
800 801 802 803 804 805 806 807 808 809 810 811 812 813 814 815 816
817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833
834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850
851 852 853 854 855 856 857 858 859 860 861 862 863 864 865 866 867
868 869 870 871 872 873 874 875 876 877 878 879 880 881 882 883 884
885 886 887 888 889 890 891 892 893 894 895 896 897 898 899 900 901
902 903 904 905 906 907 908 909 910 911 912 913 914 915 916 917 918
919 920 921 922 923 924 925 926 927 928 929 930 931 932 933 934 935
```

Inquiry Numbers 936–1122

```
936 937 938 939 940 941 942 943 944 945 946 947 948 949 950 951 952
953 954 955 956 957 958 959 960 961 962 963 964 965 966 967 968 969
970 971 972 973 974 975 976 977 978 979 980 981 982 983 984 985 986
987 988 989 990 991 992 993 994 995 996 997 998 999 1000 1001 1002 1003
1004 1005 1006 1007 1008 1009 1010 1011 1012 1013 1014 1015 1016 1017 1018 1019 1020
1021 1022 1023 1024 1025 1026 1027 1028 1029 1030 1031 1032 1033 1034 1035 1036 1037
1038 1039 1040 1041 1042 1043 1044 1045 1046 1047 1048 1049 1050 1051 1052 1053 1054
1055 1056 1057 1058 1059 1060 1061 1062 1063 1064 1065 1066 1067 1068 1069 1070 1071
1072 1073 1074 1075 1076 1077 1078 1079 1080 1081 1082 1083 1084 1085 1086 1087 1088
1089 1090 1091 1092 1093 1094 1095 1096 1097 1098 1099 1100 1101 1102 1103 1104 1105
1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122
```

Inquiry Numbers 1123–1309

```
1123 1124 1125 1126 1127 1128 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139
1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150 1151 1152 1153 1154 1155 1156
1157 1158 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1170 1171 1172 1173
1174 1175 1176 1177 1178 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1190
1191 1192 1193 1194 1195 1196 1197 1198 1199 1200 1201 1202 1203 1204 1205 1206 1207
1208 1209 1210 1211 1212 1213 1214 1215 1216 1217 1218 1219 1220 1221 1222 1223 1224
1225 1226 1227 1228 1229 1230 1231 1232 1233 1234 1235 1236 1237 1238 1239 1240 1241
1242 1243 1244 1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258
1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275
1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292
```

Inquiry Numbers 1310–1489

```
1310 1311 1312 1313 1314 1315 1316 1317 1318 1319 1320 1321 1322 1323 1324 1325 1326
1327 1328 1329 1330 1331 1332 1333 1334 1335 1336 1337 1338 1339 1340 1341 1342 1343
1344 1345 1346 1347 1348 1349 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360
1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377
1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394
1395 1396 1397 1398 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411
1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428
1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445
1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462
1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479
1480 1481 1482 1483 1484 1485 1486 1487 1488 1489
```

BYTE

FREE PRODUCT INFORMATION

For free product information from companies featured in this issue, circle the appropriate inquiry number below and mail this card today.
For quickest response, fax to 1-416-620-9790!

1. For FREE product information from individual companies, circle the corresponding inquiry numbers on the Response Card!

2. Print Your Name and Address
Answer questions "A" through "E" and mail or fax card to 1-416-620-9790.

3. Product information will be rushed to you from the selected companies!

BYTE

TNT MAILFAST
FRG/BOS/700252/MO1
P.O. Box 66
Hounslow
TW5 9RT
England

Place Postage Here

BYTE

INDEX TO ADVERTISED PRODUCTS

For FREE product information from individual advertisers, circle the corresponding inquiry numbers on the response card!

To receive information for an entire product category, circle the category number on the response card!

INDEX TO ADVERTISED PRODUCTS

EDITORIAL INDEX

For more information on any of the companies covered in articles, columns, or news stories in this issue, circle the appropriate inquiry number on the response card. Each page number refers to the first page of the article or section in which the company name appears.

IS pages appear only in the International edition. DM pages appear only in the Demographic edition.

BIX: Your Coach to the Internet!

Compatibility Testing

Tips for finding happiness in a heterogeneous environment

Are you having trouble in your personal relationships? Does it take you too long to realize you've got nothing in common with your companion? Too many dates where you discover that the person is a psycho, but you failed to read the clues? Well, friends, don't despair. Computers can help.

You can avoid disastrous dates and even find a soul mate by using this simple technique: Study how that person uses a computer. After all, computers reveal a lot about a person's temperament, philosophy of life, and mental well-being.

If, for instance, the person you're interested in uses three or four virus-scanning programs simultaneously and has devised seven levels of password protection (including some based on tokens and biometric input) just to get past the DOS prompt, you might be dealing with someone who is, shall we say, a tad insecure.

Want to know how someone feels about the whales and environmental issues? Check their desktop to see if they leave lots of applications open. If they do, chances are they are not attuned to Planet Earth. Memory is like water: There's lots of it around, but it's not unlimited. A person who wastes memory is probably careless with other limited resources, too.

What other traits can you tell from studying a person's computer habits? How about taste? (Like, if they have any.) You can bet the woman with argyle wallpaper and the Hotdog Stand desktop (bright red on brighter yellow) doesn't own many Donna Karan outfits. And the guy with a start-up screen that shows him standing next to the bloodied deer he bagged on his last hunting trip probably doesn't wear Armani.

Many things reveal a computer user's hidden personality traits. Is he or she easily impressed by size? Power? Clock speed? You'll be a disappointment until you've gotten that big promotion. Does he or she always want products with glitzy features that cost a lot? Be careful. This person is going to expect a Caribbean vacation every time a little anniversary rolls around. Computing habits provide so many clues to the real inner person.

Software-purchasing habits are also a good indicator of how a person will be in a relationship. After all, whether you're looking for a mate or buying software, you're really looking for two things: compatibility and as little retraining as possible. Some people spend a lifetime looking for the perfect one; for others, three or four will do.

Once you've found that special person or software package, you're dealing with the same set of issues. There's that excitement at the start of a relationship that's the same as exploring the features of a new spreadsheet, word processor, or object-oriented development toolkit.

It's always fun at the beginning, but some people quickly tire of the familiar. Studying a person's software upgrade habits will give you a good indication of how that person will respond when the excitement of a new relationship dies down. People who rush out and buy upgrades as soon as they're available are likely to dump you for what they perceive to be a more enticing relationship.

The irony is that most often, nothing really changes. When you upgrade to a new software package, you typically use the same features the older version had. Likewise, most people settle into the same routine they had with the previous partner; they eat at the same restaurants, go to the same movies, watch the same TV shows, etc., etc., etc.

Is loyalty important to you? If you're seeking someone who will stick it out in a relationship, find someone who hasn't upgraded their software in years. Look for someone who's still using DOS versions of their favorite programs. (I'd like to point out that this article was written with XyWrite II Plus 1.0, which was released in 1983.)

Once you've found that special someone, you can use your computer skills to make the relationship work. Consider, for example, how you'd handle an interfaith marriage. You have to make adjustments. You have to fine-tune some settings. You have to honor the other person's holidays and try to understand their perspective.

Well, that's a snap if you've worked in a mixed-operating-system or multiprotocol environment. Dealing with the most ardent religious zealot is nothing compared with trying to get a Mac disciple to use an IBM PC or an old MCI Mail user to get onto the Internet.

The bottom line is that computers reveal a great deal about a person to potential suitors. They're not called *personal* computers for nothing. ∎

Sal Salamone is a news editor who is located in BYTE's Manhattan bureau. You can reach him on the Internet or BIX at ssalamone@bix.com.

BARD MARTIN © 1995